



ONTAP 9.8 REST API reference

REST API reference

NetApp
September 12, 2025

This PDF was generated from <https://docs.netapp.com/us-en/ontap-restapi-98/index.html> on September 12, 2025. Always check docs.netapp.com for the latest.

Table of Contents

ONTAP 9.8 REST API reference	1
REST API reference	2
Getting started with the ONTAP REST API	2
Using the ONTAP REST API online documentation	2
Features for all ONTAP APIs	2
Application	24
Application overview	24
Retrieve applications	39
Create an application	88
Retrieve application components	139
Retrieve application component Snapshot copies	156
Create an application component Snapshot copy	163
Delete an application component Snapshot copy	169
Retrieve a Snapshot copy for a specific application component	172
Restore an application component Snapshot copy	178
Retrieve an application component	181
Manage application Snapshot copies	197
Delete an application and all associated data	224
Retrieve an application	228
Update application properties	271
Retrieve application templates	306
Retrieve an application template	338
Cloud	369
Cloud overview	369
Manage cloud targets	369
Cluster	413
Cluster overview	413
Manage clusters	414
Retrieve cluster chassis	510
View and manage cluster jobs	525
Retrieve capacity pool licenses	538
Manage cluster license managers	553
Manage cluster licensing	568
Manage cluster mediators	597
Cluster NTP	903
Manage cluster NTP keys	904
Manage cluster NTP servers	922
Manage cluster peers	947
Manage cluster schedules	989
Manage cluster software	1019
Manage	1087
Private manage event-remediations	1087
Event remediation and management action	1093

Perform an event management action	1097
NAS	1103
NAS overview	1103
Private protocols audit audit-log-redirect	1104
Manage NAS audit configurations	1118
Manage home directory search paths	1165
Manage CIFS services	1184
View CIFS active sessions	1270
Manage CIFS shares	1290
Manage share-level ACL	1335
Manage CIFS UNIX symlink mapping	1358
Retrieve trace results for denied or allowed events	1386
Delete trace results	1395
Retrieve trace results	1397
Retrieve security trace filter entry information	1404
Create security trace filter entries	1411
Update security trace filter entries	1416
Retrieve security trace filter entry information	1420
Update security trace filter entries	1425
View file security permissions	1428
Manage security-related accounts	5061
Manage scoped user accounts	5079
View and update audit settings	5095
Forward audit logs to syslog/splunk servers	5102
View administrative audit logs	5123
Manage data SVM account information	5132
Manage LDAP server configuration	5145
Manage NIS configuration	5169
Manage SAML service	5183
Update user account password	5200
Manage authentication keys (for admins)	5205
Manage authentication keys (end-users)	5217
Manage Azure Key Vaults	5229
Manage security certificates	5261
View and update IPsec configuration	5306
Manage IPsec policies	5311
View IPsec security associations	5336
Manage key managers	5352
Migrate SVM keys between security key managers	5416
List key servers configured in an external key manager	5420
Add primary key servers to an external key manager	5426
Delete a primary key server	5431
Retrieve key servers configured in an external key manager	5433
Update a primary key server	5438
View and update login message configuration	5442

Manage security roles	5469
View or delete a role	5485
Manage role privilege details	5494
Manage role privilege path	5503
Manage SSH server	5511
SnapLock	5522
SnapLock overview	5522
Manage SnapLock audit logs	5523
Manage the SnapLock ComplianceClock	5559
Manage event based retention operations	5567
Manage event based retention policies	5593
View key information about files and volumes	5618
Manage SnapLock file retention time	5638
Manage SnapLock storage litigations	5650
View litigation ID files	5674
Manage legal-hold operations	5678
SnapMirror	5689
SnapMirror overview	5689
Manage SnapMirror policies	5690
Manage SnapMirror relationships	5729
Manage SnapMirror relationship tranfers	5831
Storage	5864
Storage overview	5864
Retrieve or create a collection of storage aggregates	5864
Retrieve a collection of cloud stores used by an aggregate	5918
Attach an object store to an aggregate or add a second object store as a mirror	5924
Remove a cloud target from an aggregate	5930
Retrieve the cloud store for an aggregate	5933
Update a cloud store	5938
Manage storage aggregate plexes	5943
Manage storage aggregates	5963
Storage aggregate metrics	6023
Report cluster-wide storage details across different tiers	6037
Manage disks	6042
Manage storage file clones	6071
Start a file move operation	6084
Manage FlexCache volumes	6094
Manage storage FlexCache origin volumes	6132
View and create monitored files	6151
Retrieves historical performance metrics for the monitored file	6177
Manage storage ports	6186
Manage QoS policies	6203
Manage storage qtrees	6234
Storage quota	6290
Manage storage quota reports	6291

Manage storage quota policy rules	6335
Manage shelves	6393
Manage Snapshot copy policies	6421
Manage Snapshot copy policies and schedules	6450
Create a volume efficiency policy	6476
Delete a volume efficiency policy	6482
Manage storage volumes	6484
Manage files and directories	6726
Manage volume Snapshot copies	6847
Support	6887
Support overview	6887
Manage AutoSupport configuration	6887
Manage AutoSupport messages	6906
Manage cluster configuration backup information	6927
Manage configuration backups	6936
Manage EMS configuration	6959
Manage event destinations	6968
Manage an event destination instance	6982
Retrieve events	6995
Create or retrieve event filters	7002
Manage event filters	7018
Create or retrieve filter rules	7031
Manage event filter rules	7043
Manage EMS event catalog	7055
Manage cluster-wide SNMP configuration	7061
Create or retrieve SNMP traphosts	7068
Manage individual SNMP traphosts	7082
Retrieve and create SNMP users	7091
Manage individual SNMP users	7112
Legal notices	7129
Copyright	7129
Trademarks	7129
Patents	7129
Privacy policy	7129
Open source	7129

ONTAP 9.8 REST API reference

REST API reference

Getting started with the ONTAP REST API

ONTAP adds support for an expansive RESTful API. The documentation below provides information about the types of API calls available to you, as well as details about using each API endpoint. You can learn more about the ONTAP REST API and ONTAP in the ONTAP 9 Documentation Center: <http://docs.netapp.com/ontap-9/topic/com.netapp.doc.dot-rest-api/home.html>. NetApp welcomes your comments and suggestions about the ONTAP REST API and the documentation for its use.

Using the ONTAP REST API online documentation

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

Features for all ONTAP APIs

Getting started with the ONTAP REST API

Overview

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

- REST API URLs identify the resources that you'll be working with, including clusters, SVMs, and storage.
- REST APIs use HTTP methods GET, POST, PATCH, DELETE, and OPTIONS to indicate their actions.
- REST APIs return common HTTP status codes to indicate the results of each call. Additional error details can be included in the results body.
- REST APIs request and response bodies are encoded using JSON.
- REST APIs support hyperlinking among resources using the Content-Type "application/hal+json".
- GET calls on collections usually return only name and UUID by default. If you want to retrieve additional properties, you need to specify them using the "fields" query parameter.
- ONTAP supports query-based DELETE or PATCH for all collection endpoints.
If you're already familiar with the ONTAPI API (also known as ZAPI), there are some similarities between ONTAP REST APIs and ONTAPI. For example:
 - Both support the same transport and security mechanisms.
 - Both paginate results based on either number of seconds or number of records.
 - Both support filtering the returned records based on property values.
 - Both support limiting the returned properties.However, there are important differences between REST APIs and the ONTAP CLI and ONTAPI that you should understand as well:
 - In many cases, ONTAP REST APIs use different names for fields and features.
 - REST APIs do not expose infrequently used CLI parameters.
 - REST APIs do not treat the cluster or nodes as an SVM (aka Vserver).
 - REST GET APIs support specifying a maximum time before paginating results. However, the default time is 15 seconds for REST (instead of 90 seconds for ONTAPI).

- REST APIs are generally ordered by UUID or ID, so a rename operation using the PATCH method doesn't change the path keys.
- REST APIs use one or more of the following properties to identify a resource: "name", "uuid", "id".
- REST APIs often execute the equivalent of multiple CLI commands in a single request.
- REST API properties use underscores instead of hyphens between words.
- REST API dates are always in ISO-8601 format.
- REST API comparisons between enum values (for <, >, ranges, and `order_by`) are done alphabetically. (In CLI and ONTAPI, enum comparisons are done based on an internal value for the enum.)
- REST API field '<' queries exclude records where the specified field is not set. You can add "\null" (eg: `limit=<10\null`) to also return records where the specified field is not set.

HAL linking

Hypertext Application Language (HAL)

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

Example

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

Query parameters

Overview

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

```

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

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

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

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

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

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

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

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

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

```

Query-based PATCH and DELETE

Overview

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

Example

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

Record filtering

Overview

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

Filtering records with single field queries

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

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

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

Wildcard: *

```
abc*
abc*xyz

***xyz
```

Comparison: < > <= >=

```
<10
>=joe
```

Range: ..

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

Negation: !

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

Any of a list: |

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

Escaping: {} and ""

The special query characters above can be treated literally, with no special meaning, by enclosing the value in either double quotes or curly braces.

```
"joe*"
{a|b}
```

Filtering records with cross-field queries

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

The fields to be queried across are specified in the "query_fields=" parameter. This should be a comma-delimited list of fields, or simply * to search across all fields.

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

Structure of the query

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

The * character is used to indicate wildcard character matching. * matches 0 or more of any character. For a query of "foo*bar", matches will include "foo123abcbar", "foobar", and "foo__123abcbar".

To search for any match among several possible patterns, the values may be ORed together with the '|' character. For example, to search for "foo" OR "bar", pass in "query=foo|bar". This may be extended to an arbitrary number of values, such as "query=foo|bar|baz".

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

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

Examples

The following data is used for the examples below:

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

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

Response:

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

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

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

Response:

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

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

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

Response:

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

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

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

Response:

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

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

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

Response:

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

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

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

Response:

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

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

Cross-Field Query Errors

ONTAP Error Response Codes

Error Code	Description	HTTP Code
262272	The specified query contains an unmatched quote.	400
262273	Both 'query_fields' and 'query' must be specified if either one is specified.	400

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

Requesting specific fields

Overview

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

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

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

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

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

Objects with fields

When an API contains fields that are objects, an entire object can be specified to return every field within the object. Individual fields within the object can be specified using dotted notation, as demonstrated below. Occasionally, some APIs might not support specifying fields nested within certain objects.



Dotted notation for arrays does not include array indices.

Examples

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

Example fields query:

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

Records and pagination

Records

Several query parameters control the return of records.

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

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

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

The `return_timeout` parameter specifies the number of seconds the cluster spends performing an operation before returning. The allowed range is 0 to 120 seconds. If the timeout is reached, GET calls return the records collected along with a pagination link. Other methods return and complete asynchronously. See [Non-blocking-operations](#) for more details.

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



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

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

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

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

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

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

Pagination

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

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

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

Count only

The response to collection operations includes a `num_records` field. By passing `return_records=false` with a GET call, you can retrieve the number of records without returning the records themselves. However, if either the `return_timeout` or `max_records` threshold is reached, an incomplete or partial number of records is returned and the "next" link must be called to retrieve additional record counts. All the partial counts

must be added together to calculate the total count.

Record sorting

Overview

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

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

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

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

Important Notes:

- When you use the `order_by` parameter, the `return_timeout` might be exceeded because the collection is sorted before it is returned.
- Using `order_by` on either a property of type array, or a nested property within an array (not including the records array), returns the records in an unspecified order.

Examples

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

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

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

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

Response body

Overview

Every API call returns a top-level JSON object. These JSON objects includes GET calls that contain an array

of records. This nesting technique allows metadata about the resource or resource collection to be returned as well as each resource instance.

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

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

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

Custom response bodies

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

Error objects

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

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

ONTAP Error Response Codes

Error Code	Description	HTTP Code
1	An entry with the same identifiers already exists.	409
2	A field has an invalid value, is missing, or an extra field was provided.	400
3	The operation is not supported.	405
4	An entry with the specified identifiers was not found.	404
6	Permission denied.	403
8	Resource in use.	409
65541	RPC timed out.	500
65552	Generic RPC failure.	500
65562	Internal RPC error	500
262177	Missing value.	400
262179	Unexpected argument. Argument shown in error message body.	400
262185	Invalid value with value in the body of the error.	400
262186	A field is used in an invalid context with another field, as shown in error message body.	400
262188	A field was specified twice. Location of assignments shown in error message body.	400
262197	Invalid value provided for field. Value and field shown in error message body.	400
262198	A request body is not allowed on GET, HEAD, and DELETE.	400
262199	Invalid JSON with error location provided in body of the error.	400
262200	Invalid JSON range, with range provided in the body of the error.	400
262201	Invalid JSON due to unknown formatting issue.	400
262245	Invalid value with reason provided in body of the error.	400
262247	Invalid value for a field, with value and field in body of the error.	400

Error Code	Description	HTTP Code
262248	A value is missing assignment operator.	400
262255	An array was found in the JSON when it was not expected.	400
393271	A node is out of quorum. Body of error message identifies node.	500
39387137	A provided URL is invalid.	400

Synchronous and asynchronous operations

Overview

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

API response

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

The Location header

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

Tracking non-blocking operations

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

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

See [GET /cluster/jobs](#)

HTTP status codes

Overview

The following supported HTTP status codes are returned by ONTAP:

- 200 OK: Returned for success when not creating a new object
- 201 Created: Returned for success after the creation of an object
- 202 Accepted: Returned when a job has been successfully started, but the operation is not complete
- 400 Bad Request: Returned if the input could not be parsed
- 401 Unauthorized: Returned if user authentication failed
- 403 Forbidden: Returned for authorization (RBAC) errors
- 404 Not Found: Returned when the specified resource does not exist
- 405 Method Not Allowed: Returned when the specified resource does not support the method (for example, POST or DELETE calls)
- 409 Conflict: Returned when there is a conflict with a different object that must be created, modified, or deleted before this operation can succeed
- 500 Internal Error: Returned for all other internal error codes

HTTP methods

Overview

The ONTAP REST API supports the following HTTP methods:

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

Size properties

Overview

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

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

Depending on the development language-specific code generation, the API typically also requires an integer value in bytes for POST and PATCH input as well.

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

Suffix	Definition
KB	kilobytes (1024 bytes, aka kibibytes)
MB	megabytes (KB x 1024, aka mebibytes)

Suffix	Definition
GB	gigabytes (MB x 1024, aka gibibytes)
TB	terabytes (GB x 1024, aka tebibytes)
PB	petabytes (TB x 1024, aka pebibytes)

SVM tunneling

Overview

SVM tunneling allows for the scoping of REST APIs to any SVM from the cluster admin SVM interface. The HTTP headers "X-Dot-SVM-Name" and/or "X-Dot-SVM-UUID" are an alternative to supplying svm.name and/or svm.uuid in the request query or body. This allows for setting a context for an HTTP connection and reusing it for multiple calls. The cluster management interface or node management interface can be used instead of the desired SVM's interface.

Examples

Creates a new volume on SVM "vs0":

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

Retrieves all volumes on SVM "vs0":

```

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

```

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

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

Retrieves all IP interfaces on SVM "vs3":

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

Using the private CLI passthrough with the ONTAP REST API

REST API access to CLI commands

To help CLI and ONTAP users transition to the ONTAP REST API, ONTAP provides a private REST API endpoint that can be used to access any CLI command. Usage of this API call is recorded and returned in the

AutoSupport data collection so that NetApp can identify usability and functionality improvements in the REST API for future releases. There is no per-API documentation for the REST API access for each CLI command. Unlike the documented REST APIs, the API paths and properties for the CLI passthrough correspond very closely to the CLI. There are several rules that govern all the differences between a CLI command and the REST API mirroring the CLI command.

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

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

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

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

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

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

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

All CLI parameters are supported in the CLI-based REST APIs. However, REST converts hyphens (-) in CLI

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

Examples

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

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

GET size and percent-used for all volumes:

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

GET size and percent-used for a specific volume:

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

POST a new volume with all required attributes:

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

Attempt to DELETE an online volume:

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

PATCH a volume to become offline:

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

DELETE the offline volume:

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



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

DELETE an offline volume without waiting:

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

CLI message output

As shown in the previous example, any non-field and non-error based output that would have appeared in the

CLI is returned in a top-level `cli_output` attribute in the response body. This does not contain normal CLI headers or field values. It only displays messages that were printed to the CLI.

HTTP status codes

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

Security

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

Location of CLI fields for CLI-based REST APIs:

- POST APIs: All CLI fields must be provided in the request body.
- GET APIs: All desired CLI fields (except keys) must be specified in the `fields` parameter. The non-key fields returned via the CLI will not be returned if not requested. The client can also provide a query for any field.
- PATCH APIs: The client can provide a query for any field, but at least one field must have a query. To modify only a single record, all CLI keys must contain an exact query. All new values for the object must be provided in the request body.
- DELETE APIs: The client can provide a query for any field, but at least one field must have a query. To delete only a single record, all CLI keys must contain an exact query. Non-attribute inputs (such as `force`) must be provided in the query portion of the URI.

Application

Application overview

Overview

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

APIs

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

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

Template

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

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

Application

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

The application object includes the following three sections:

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

Component

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

The component object includes the following details for an application:

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

Snapshot copy

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

Example

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

1) Discover the templates

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

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

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

# The query:
name=oracle*

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

# The response:
{
  "records": [
    {
      "name": "oracle_on_nfs", "description": "Oracle using NFS."
    },
    {
      "name": "oracle_on_san", "description": "Oracle using SAN."
    },
    {
      "name": "oracle_rac_on_nfs", "description": "Oracle RAC using NFS."
    },
    {
      "name": "oracle_rac_on_san", "description": "Oracle RAC using SAN."
    }
  ],
  "num_records": 4
}
```

2) Create an application

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



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

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

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

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

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

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

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

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

# The body:
{
  "name": "my_ora_app",
  "svm": {
    "name": "svm1"
  },
  "oracle_on_nfs": {
    "db": {
      "size": "2GB",
      "storage_service": {
        "name": "value"
      },
    },
  },
}
```

```

    "nfs_access": [
        {
            "access": "rw",
            "host": "0.0.0.0/0"
        }
    ]
},
"redo_log": {
    "size": "1GB"
},
"ora_home": {
    "size": "1GB"
}
}
}

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

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

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

```

3) Wait for the application to be created

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

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

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

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

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

4) Retrieve the new application

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



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

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

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

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

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

5) Discover how to access the application

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

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

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

# The query:
fields=name

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

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

6) Update the application

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

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

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

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

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

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

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

7) Manage Snapshot copies

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

Take a Snapshot copy manually:

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

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

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

# The response:
{}

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

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

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

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

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

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

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

Smart containers

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

- NAS - FlexVolume, FlexGroup, and FlexCache objects
- SAN - LUNs
- NVME - Namespaces

The benefits of creating a Smart Container are as follows:

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

Smart containers are similar to generic enterprise applications (NAS, SAN, NVME), but with certain restrictions. Smart containers are restricted to 1 application-component. Any post-provisioning data management operations on smart containers must be performed via PATCH operations corresponding to the object created. However, the POST, GET and DELETE operations that exist for applications will also operate for smart containers.

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

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

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

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

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

Example

The following example outlines the APIs necessary to create a smart container. This is an addendum to the example provided on how to create an application.

```

# The API:
/api/application/applications

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

# The body:
{
  "name": "my_container",
  "svm": {
    "name": "vs1"
  },
  "template": {
    "name": "nas"
  },
  "smart_container": "true"
  "nas": {
    "application_components": [
      {
        "share_count": "1",
        "name": "myVolume",
        "storage_service": {
          "name": "value"
        },
        "total_size": "100mb"
      }
    ]
  }
}

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

# The response:
{
  "job": {
    "uuid": "5440db05-77f0-11e9-a5a0-005056bba32f",
    "_links": {
      "self": {

```

```
      "href": "/api/cluster/jobs/5440db05-77f0-11e9-a5a0-005056bba32f"
    }
  }
}
```

The response header:

```
date: Tue, 23 May 2019 16:38:18 GMT
server: libzapid-httpd
content-type: application/hal+json
location: /api/application/applications/5440db05-77f0-11e9-a5a0-005056bba32f
cache-control: no-cache,no-store,must-revalidate
connection: Keep-Alive
keep-alive: timeout=5, max=100
content-length: 203
```

Smart Container Properties

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

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

Application API limitations

Template versus generic

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

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

ONTAP feature support

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

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

Retrieve applications

GET `/application/applications`

Introduced In: 9.6

Retrieves applications.

Expensive properties

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

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

Query examples

Numerous queries are available for classifying and sorting applications:

1. Return a list of applications sorted by name.

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

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

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

3. Return a list of all SQL applications.

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

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


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

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

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

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

Learn more

- [DOC /application](#)

Parameters

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

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

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

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

Name	Type	In	Required	Description
statistics.component s.space.reserved_un used	string	query	False	Filter by statistics.component s.space.reserved_u nused
statistics.component s.space.available	string	query	False	Filter by statistics.component s.space.available
statistics.component s.space.savings	string	query	False	Filter by statistics.component s.space.savings
statistics.component s.iops.total	string	query	False	Filter by statistics.component s.iops.total
statistics.component s.iops.per_tb	string	query	False	Filter by statistics.component s.iops.per_tb
statistics.component s.snapshot.reserve	string	query	False	Filter by statistics.component s.snapshot.reserve
statistics.component s.snapshot.used	string	query	False	Filter by statistics.component s.snapshot.used
statistics.component s.latency.raw	string	query	False	Filter by statistics.component s.latency.raw
statistics.component s.latency.average	string	query	False	Filter by statistics.component s.latency.average
statistics.component s.statistics_incomple te	string	query	False	Filter by statistics.component s.statistics_incomple te
statistics.component s.shared_storage_p ool	string	query	False	Filter by statistics.component s.shared_storage_p ool

Name	Type	In	Required	Description
smart_container	string	query	False	Filter by smart_container <ul style="list-style-type: none"> Introduced in: 9.7
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none"> Default value: 1 Max value: 120 Min value: 0
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

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

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "snapshots": {
        "href": "/api/resourcelink"
      }
    },
    "cg_on_san": {
      "cgs": [
        {
          "lun_containers": [
            {
              "luns": [
                {
                  "igroups": [
                    {
                      "name": "string"
                    }
                  ],
                  "name": "string",
                  "os_type": "string"
                }
              ],
              "name": "string"
            }
          ],
          "name": "string",
          "protection_type": {
            "local_policy": "string",
            "remote_rpo": "string"
          },
          "qos": {
```



```

        "policy": {
            "name": "string",
            "uuid": "string"
        }
    },
    "storage_service": {
        "name": "string"
    },
    "tiering": {
        "policy": "string"
    }
}
]
},
"creation_timestamp": "string",
"delete_data": null,
"generation": 0,
"maxdata_on_san": {
    "app_type": "string",
    "application_components": [
        {}
    ],
    "metadata": [
        {
            "key": "string",
            "value": "string"
        }
    ],
    "ocsm_url": "string",
    "os_type": "string"
},
"name": "string",
"nas": {
    "application_components": [
        {}
    ],
    "cifs_access": [
        {
            "access": "string",
            "user_or_group": "string"
        }
    ],
    "nfs_access": [
        {
            "access": "string",
            "host": "string"
        }
    ]
}

```

```

    }
  ],
  "protection_type": {
    "local_policy": "string",
    "local_rpo": "string",
    "remote_rpo": "string"
  }
},
"nvme": {
  "components": [
    {
      "name": "string",
      "os_type": "string",
      "performance": {
        "storage_service": {
          "name": "string"
        }
      },
    },
    "qos": {
      "policy": {
        "name": "string",
        "uuid": "string"
      }
    },
    "subsystem": {
      "hosts": [
        {
          "nqn": "string"
        }
      ],
      "name": "string",
      "os_type": "string",
      "uuid": "string"
    },
    "tiering": {
      "policy": "string"
    }
  ]
},
"os_type": "string",
"rpo": {
  "local": {
    "name": "string",
    "policy": "string"
  }
}

```

```

},
"protection_granularity": "string",
"rpo": {
  "components": [
    {
      "name": "string",
      "rpo": {
        "local": {
          "description": "string",
          "name": "string"
        },
        "remote": {
          "description": "string",
          "name": "string"
        }
      },
      "uuid": "string"
    }
  ],
  "local": {
    "description": "string",
    "name": "string"
  },
  "remote": {
    "description": "string",
    "name": "string"
  }
},
"s3_bucket": {
  "application_components": [
    {}
  ]
},
"san": {
  "application_components": [
    {
      "igroup_name": "string",
      "name": "string",
      "os_type": "string",
      "qos": {
        "policy": {
          "name": "string",
          "uuid": "string"
        }
      },
      "storage_service": {

```

```

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

```

```

    }
  ],
  "iops": {
    "per_tb": 0,
    "total": 0
  },
  "latency": {
    "average": 0,
    "raw": 0
  },
  "snapshot": {
    "reserve": 0,
    "used": 0
  },
  "space": {
    "available": 0,
    "logical_used": 0,
    "provisioned": 0,
    "reserved_unused": 0,
    "savings": 0,
    "used": 0,
    "used_excluding_reserves": 0,
    "used_percent": 0
  }
},
"svm": {
  "name": "string",
  "uuid": "string"
},
"template": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "string",
  "protocol": "string",
  "version": 0
},
"uuid": "string"
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	
snapshots	href	

igroups

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

zapp_cg_on_san_cgs_lun_containers_luns

The list of LUNs to be created.

Name	Type	Description
count	integer	The number of LUNs.
igroups	array[igroups]	
name	string	The name/prefix of LUNs.
os_type	string	The name of the host OS.

zapp_cg_on_san_cgs_lun_containers

LUN containers.

Name	Type	Description
luns	array[zapp_cg_on_san_cgs_lun_containers_luns]	The list of LUNs to be created.
name	string	The name or prefix of container (Volume). If the container exists, it becomes a part of the CG. If not, a new container is created.

protection_type

Name	Type	Description
local_policy	string	The Snapshot copy policy to apply to a CG.
remote_rpo	string	The remote RPO of the CG.

policy

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

qos

Name	Type	Description
policy	policy	

storage_service

Name	Type	Description
name	string	The storage service of the CG.

object_stores

zapp_cg_on_san_cgs_tiering

tiering

Name	Type	Description
policy	string	The storage tiering type of the CG.

zapp_cg_on_san_cgs

The list of CGs to be created.

Name	Type	Description
lun_containers	array[zapp_cg_on_san_cgs_lun_containers]	LUN containers.
name	string	The name of the consistency group.
protection_type	protection_type	
qos	qos	
storage_service	storage_service	
tiering	zapp_cg_on_san_cgs_tiering	tiering

zapp_cg_on_san

A consistency group (CG) application using SAN.

Name	Type	Description
cgs	array[zapp_cg_on_san_cgs]	The list of CGs to be created.

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application component.
value	string	Value associated with the key.

protection_type

Name	Type	Description
local_rpo	string	The local rpo of the application component.

Name	Type	Description
remote_rpo	string	The remote rpo of the application component.

storage_service

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

maxdata_on_san_application_components_tiering

tiering

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

maxdata_on_san_application_components

The list of application components to be created.

Name	Type	Description
file_system	string	Defines the type of file system that will be installed on this application component.
host_management_url	string	The host management URL for this application component.
host_name	string	FQDN of the L2 host that contains the hot tier of this application component.
igroup_name	string	The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup.

Name	Type	Description
lun_count	integer	The number of LUNs in the application component.
metadata	array[metadata]	
name	string	The name of the application component.
protection_type	protection_type	
storage_service	storage_service	
tiering	maxdata_on_san_application_components_tiering	tiering
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]}

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application.
value	string	Value associated with the key.

maxdata_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

maxdata_on_san

MAX Data application using SAN.

Name	Type	Description
app_type	string	Type of the application that is being deployed on the L2.
application_components	array[maxdata_on_san_application_components]	The list of application components to be created.
metadata	array[metadata]	

Name	Type	Description
ocsm_url	string	The OnCommand System Manager URL for this application.
os_type	string	The name of the host OS running the application.

storage_service

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

dataset

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

mongo_db_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

protection_type

Name	Type	Description
local_rpo	string	The local RPO of the application.

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

secondary_igroups

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

mongo_db_on_san

MongoDB using SAN.

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

export_policy

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

component

Name	Type	Description
name	string	Name of the source component.

svm

Name	Type	Description
name	string	Name of the source SVM.

origin

Name	Type	Description
component	component	
svm	svm	

flexcache

Name	Type	Description
origin	origin	

nas_application_components_tiering

application-components.tiering

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

application_components

Name	Type	Description
export_policy	export_policy	
flexcache	flexcache	
name	string	The name of the application component.
qos	qos	
scale_out	boolean	Denotes a Flexgroup.
share_count	integer	The number of shares in the application component.
storage_service	storage_service	
tiering	nas_application_components_tiering	application-components.tiering
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]}

app_cifs_access

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

access.

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

app_nfs_access

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

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

protection_type

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

nas

A generic NAS application.

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

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

performance

Name	Type	Description
storage_service	storage_service	

hosts

Name	Type	Description
nqn	string	The host NQN.

zapp_nvme_components_subsystem

components.subsystem

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

zapp_nvme_components_tiering

application-components.tiering

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

components

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

local

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

rpo

Name	Type	Description
local	local	

zapp_nvme

An NVME application.

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

storage_service

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

archive_log

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

db

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

storage_service

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

ora_home

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

storage_service

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

redo_log

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

oracle_on_nfs

Oracle using NFS.

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

oracle_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

oracle_on_san

Oracle using SAN.

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

storage_service

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

grid_binary

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

storage_service

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

oracle_crs

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

oracle_rac_on_nfs

Oracle RAC using NFS.

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

db_sids

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

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

oracle_rac_on_san

Oracle RAC using SAN.

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

local

Name	Type	Description
description	string	A detailed description of the local RPO. This will include details about the Snapshot copy schedule.

Name	Type	Description
name	string	The local RPO of the component. This indicates how often component Snapshot copies are automatically created.

remote

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

rpo

Name	Type	Description
local	local	
remote	remote	

components

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

local

Name	Type	Description
description	string	A detailed description of the local RPO. This will include details about the Snapshot copy schedule.

Name	Type	Description
name	string	The local RPO of the application. This indicates how often application Snapshot copies are automatically created.

remote

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

rpo

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

zapp_s3_bucket_application_components_access_policies_conditions

conditions

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

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

Name	Type	Description
actions	array[string]	
conditions	array[zapp_s3_bucket_application_components_access_policies_conditions]	conditions.
effect	string	Allow or Deny Access.
principal	array[string]	
resources	array[string]	
sid	string	Statement Identifier.

exclude_aggregates

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

zapp_s3_bucket_application_components

The list of application components to be created.

Name	Type	Description
access_policies	array[zapp_s3_bucket_application_components_access_policies]	The list of S3 objectstore policies to be created.
capacity_tier	boolean	Prefer lower latency storage under similar media costs.
comment	string	Object Store Server Bucket Description Usage: <(size 1..256)>
exclude_aggregates	array[exclude_aggregates]	
name	string	The name of the application component.
qos	qos	

Name	Type	Description
size	integer	The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB MB GB TB PB]}
storage_service	storage_service	
uuid	string	Object Store Server Bucket UUID Usage: <UUID>

zapp_s3_bucket

A generic S3 bucket application.

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

san_application_components_tiering

application-components.tiering

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

application_components

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

Name	Type	Description
os_type	string	The name of the host OS running the application.
qos	qos	
storage_service	storage_service	
tiering	san_application_components_tiering	application-components.tiering
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]}

san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

san

A generic SAN application.

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

storage_service

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

db

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

storage_service

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

log

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

sql_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

storage_service

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

temp_db

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

sql_on_san

Microsoft SQL using SAN.

Name	Type	Description
db	db	

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

access

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

sql_on_smb

Microsoft SQL using SMB.

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

iops

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

latency

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

snapshot

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

space

Name	Type	Description
available	integer	<p>The available amount of space left in the application component. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.</p> <ul style="list-style-type: none"> • readOnly: 1

Name	Type	Description
logical_used	integer	The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used.
provisioned	integer	The originally requested amount of space that was provisioned for the application component.
reserved_unused	integer	The amount of space reserved for system features such as Snapshot copies that has not yet been used.
savings	integer	The amount of space saved by all enabled space saving features.
used	integer	The amount of space that is currently being used by the application component. Note that this includes any space reserved by the system for features such as Snapshot copies.
used_excluding_reserves	integer	The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies.
used_percent	integer	The percentage of the originally provisioned space that is currently being used by the application component.

storage_service

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

Name	Type	Description
uuid	string	The storage service UUID.

components

Name	Type	Description
iops	iops	
latency	latency	
name	string	Component Name.
shared_storage_pool	boolean	An application component is considered to use a shared storage pool if storage elements for other components reside on the same aggregate as storage elements for this component.
snapshot	snapshot	
space	space	
statistics_incomplete	boolean	If not all storage elements of the application component are currently available, the returned statistics might only include data from those elements that were available.
storage_service	storage_service	
uuid	string	Component UUID.

iops

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

latency

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

space

Name	Type	Description
available	integer	<p>The available amount of space left in the application. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.</p> <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6
logical_used	integer	The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used.
provisioned	integer	The originally requested amount of space that was provisioned for the application.
reserved_unused	integer	The amount of space reserved for system features such as Snapshot copies that has not yet been used.
savings	integer	The amount of space saved by all enabled space saving features.

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

statistics

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

svm

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

self_link

Name	Type	Description
self	href	

template

Name	Type	Description
_links	self_link	
name	string	The name of the template that was used to provision this application.
protocol	string	The protocol access of the template that was used to provision this application.
version	integer	<p>The version of the template that was used to provision this application. The template version changes only if the layout of the application changes over time. For example, redo logs in Oracle RAC templates were updated and provisioned differently in DATA ONTAP 9.3.0 compared to prior releases, so the version number was increased. If layouts change in the future, the changes will be documented along with the corresponding version numbers.</p> <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6

storage_service

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

desktops

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

hyper_v_access

Name	Type	Description
service_account	string	Hyper-V service account.

vdi_on_nas

A VDI application using NAS.

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

vdi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

vdi_on_san

A VDI application using SAN.

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

storage_service

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

datastore

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

vsi_on_nas

A VSI application using NAS.

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

vsi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

vsi_on_san

A VSI application using SAN.

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

application

Applications

Name	Type	Description
_links	_links	
cg_on_san	zapp_cg_on_san	A consistency group (CG) application using SAN.
creation_timestamp	string	The time when the application was created.

Name	Type	Description
generation	integer	The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented.
maxdata_on_san	maxdata_on_san	MAX Data application using SAN.
mongo_db_on_san	mongo_db_on_san	MongoDB using SAN.
name	string	Application Name. This field is user supplied when the application is created.
nas	nas	A generic NAS application.
nvme	zapp_nvme	An NVME application.
oracle_on_nfs	oracle_on_nfs	Oracle using NFS.
oracle_on_san	oracle_on_san	Oracle using SAN.
oracle_rac_on_nfs	oracle_rac_on_nfs	Oracle RAC using NFS.
oracle_rac_on_san	oracle_rac_on_san	Oracle RAC using SAN.
protection_granularity	string	Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components.
rpo	rpo	
s3_bucket	zapp_s3_bucket	A generic S3 bucket application.

Name	Type	Description
san	san	A generic SAN application.
smart_container	boolean	Identifies if this is a smart container or not.
sql_on_san	sql_on_san	Microsoft SQL using SAN.
sql_on_smb	sql_on_smb	Microsoft SQL using SMB.
state	string	The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported.
statistics	statistics	
svm	svm	
template	template	
uuid	string	Application UUID. This field is generated when the application is created.
vdi_on_nas	vdi_on_nas	A VDI application using NAS.
vdi_on_san	vdi_on_san	A VDI application using SAN.
vsi_on_nas	vsi_on_nas	A VSI application using NAS.
vsi_on_san	vsi_on_san	A VSI application using SAN.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Create an application

POST /application/applications

Introduced In: 9.6

Creates an application.

Template properties

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

Required properties

- `svm.uuid` or `svm.name` - The existing SVM in which to create the application.
- `name` - The name for the application.
- `<template>` - Properties for one template must be provided. In general, the following properties are required, however the naming of these may vary slightly from template to template.
 - `name` - The generic templates require names for the components of the application. Other templates name the components automatically.
 - `size` - This generally refers to the size of an application component, which may be spread across multiple underlying storage objects (volumes, LUNs, etc...).
 - One of the following must be specified:
- `nfs_access` or an identifier (name or id) of an existing `export-policy`.
- `cifs_access`
- `igroup_name`
 - `os_type` - All SAN applications require an `os_type` to be specified in some way. Some templates refer to this as the `hypervisor`.

Recommended optional properties

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

Default property values

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

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

Optional components

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

POST body examples

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

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

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

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

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



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

```

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

```

Learn more

- [DOC /application](#)
- [Asynchronous operations](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
_links	_links	
cg_on_san	zapp_cg_on_san	A consistency group (CG) application using SAN.

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

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

Name	Type	Description
vsi_on_san	vsi_on_san	A VSI application using SAN.

Example request

```
{
  "cg_on_san": {
    "cgs": [
      {
        "lun_containers": [
          {
            "luns": [
              {
                "igroups": [
                  {
                    "name": "string"
                  }
                ],
                "name": "string",
                "os_type": "string"
              }
            ],
            "name": "string"
          }
        ],
        "name": "string",
        "protection_type": {
          "local_policy": "string",
          "remote_rpo": "string"
        },
        "qos": {
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "storage_service": {
          "name": "string"
        },
        "tiering": {
          "control": "string",
          "object_stores": [
            {
              "name": "string"
            }
          ],
          "policy": "string"
        }
      }
    ]
  }
}
```

```

    ],
    "creation_timestamp": "string",
    "generation": 0,
    "maxdata_on_san": {
        "app_type": "string",
        "application_components": [
            {}
        ],
    },
    "metadata": [
        {
            "key": "string",
            "value": "string"
        }
    ],
    "new_igroups": [
        {
            "initiators": [
                "string"
            ],
            "name": "string",
            "os_type": "string",
            "protocol": "string"
        }
    ],
    "ocsm_url": "string",
    "os_type": "string"
},
"name": "string",
"nas": {
    "application_components": [
        {}
    ],
    "cifs_access": [
        {
            "access": "string",
            "user_or_group": "string"
        }
    ],
    "nfs_access": [
        {
            "access": "string",
            "host": "string"
        }
    ],
    "protection_type": {

```

```

    "local_policy": "string",
    "local_rpo": "string",
    "remote_rpo": "string"
  }
},
"nvme": {
  "components": [
    {
      "name": "string",
      "os_type": "string",
      "performance": {
        "storage_service": {
          "name": "string"
        }
      },
      "qos": {
        "policy": {
          "name": "string",
          "uuid": "string"
        }
      },
      "subsystem": {
        "hosts": [
          {
            "nqn": "string"
          }
        ],
        "name": "string",
        "os_type": "string",
        "uuid": "string"
      },
      "tiering": {
        "control": "string",
        "object_stores": [
          {
            "name": "string"
          }
        ],
        "policy": "string"
      }
    }
  ],
  "os_type": "string",
  "rpo": {
    "local": {
      "name": "string",

```

```

        "policy": "string"
    }
}
},
"protection_granularity": "string",
"rpo": {
    "components": [
        {
            "name": "string",
            "rpo": {
                "local": {
                    "description": "string",
                    "name": "string"
                },
                "remote": {
                    "description": "string",
                    "name": "string"
                }
            },
            "uuid": "string"
        }
    ],
    "local": {
        "description": "string",
        "name": "string"
    },
    "remote": {
        "description": "string",
        "name": "string"
    }
},
"s3_bucket": {
    "application_components": [
        {}
    ]
},
"san": {
    "application_components": [
        {
            "igroup_name": "string",
            "name": "string",
            "os_type": "string",
            "qos": {
                "policy": {
                    "name": "string",
                    "uuid": "string"
                }
            }
        }
    ]
}

```

```

    }
  },
  "storage_service": {
    "name": "string"
  },
  "tiering": {
    "control": "string",
    "object_stores": [
      {
        "name": "string"
      }
    ],
    "policy": "string"
  }
},
"new_igroups": [
  {
    "initiators": [
      "string"
    ],
    "name": "string",
    "os_type": "string",
    "protocol": "string"
  }
],
"os_type": "string",
"protection_type": {
  "local_policy": "string",
  "local_rpo": "string",
  "remote_rpo": "string"
},
},
"state": "string",
"statistics": {
  "components": [
    {
      "iops": {
        "per_tb": 0,
        "total": 0
      },
      "latency": {
        "average": 0,
        "raw": 0
      },
      "name": "string",

```

```

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

```

```
    "uuid": "string"
  },
  "template": {
    "name": "string",
    "protocol": "string",
    "version": 0
  },
  "uuid": "string"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

igroups

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

zapp_cg_on_san_cgs_lun_containers_luns

The list of LUNs to be created.

Name	Type	Description
count	integer	The number of LUNs.
igroups	array[igroups]	
name	string	The name/prefix of LUNs.
os_type	string	The name of the host OS.

zapp_cg_on_san_cgs_lun_containers

LUN containers.

Name	Type	Description
luns	array[zapp_cg_on_san_cgs_lun_containers_luns]	The list of LUNs to be created.
name	string	The name or prefix of container (Volume). If the container exists, it becomes a part of the CG. If not, a new container is created.

protection_type

Name	Type	Description
local_policy	string	The Snapshot copy policy to apply to a CG.
remote_rpo	string	The remote RPO of the CG.

policy

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

qos

Name	Type	Description
policy	policy	

storage_service

Name	Type	Description
name	string	The storage service of the CG.

object_stores

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

zapp_cg_on_san_cgs_tiering

tiering

Name	Type	Description
control	string	The storage tiering placement rules for the CG.
object_stores	array[object_stores]	

Name	Type	Description
policy	string	The storage tiering type of the CG.

zapp_cg_on_san_cgs

The list of CGs to be created.

Name	Type	Description
lun_containers	array[zapp_cg_on_san_cgs_lun_containers]	LUN containers.
name	string	The name of the consistency group.
protection_type	protection_type	
qos	qos	
storage_service	storage_service	
tiering	zapp_cg_on_san_cgs_tiering	tiering

zapp_cg_on_san

A consistency group (CG) application using SAN.

Name	Type	Description
cgs	array[zapp_cg_on_san_cgs]	The list of CGs to be created.

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application component.
value	string	Value associated with the key.

protection_type

Name	Type	Description
local_rpo	string	The local rpo of the application component.

Name	Type	Description
remote_rpo	string	The remote rpo of the application component.

storage_service

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

maxdata_on_san_application_components_tiering

tiering

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

maxdata_on_san_application_components

The list of application components to be created.

Name	Type	Description
file_system	string	Defines the type of file system that will be installed on this application component.
host_management_url	string	The host management URL for this application component.
host_name	string	FQDN of the L2 host that contains the hot tier of this application component.

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

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application.
value	string	Value associated with the key.

maxdata_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	
name	string	The name of the new initiator group.

Name	Type	Description
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application.
protocol	string	The protocol of the new initiator group.

maxdata_on_san

MAX Data application using SAN.

Name	Type	Description
app_type	string	Type of the application that is being deployed on the L2.
application_components	array[maxdata_on_san_application_components]	The list of application components to be created.
metadata	array[metadata]	
new_igroups	array[maxdata_on_san_new_igroups]	The list of initiator groups to create.
ocsm_url	string	The OnCommand System Manager URL for this application.
os_type	string	The name of the host OS running the application.

storage_service

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

dataset

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

mongo_db_on_san_new_igroups

The list of initiator groups to create.

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

protection_type

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

secondary_igroups

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

mongo_db_on_san

MongoDB using SAN.

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

export_policy

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

component

Name	Type	Description
name	string	Name of the source component.

svm

Name	Type	Description
name	string	Name of the source SVM.

origin

Name	Type	Description
component	component	
svm	svm	

flexcache

Name	Type	Description
origin	origin	

nas_application_components_tiering

application-components.tiering

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

application_components

Name	Type	Description
export_policy	export_policy	
flexcache	flexcache	
name	string	The name of the application component.
qos	qos	
scale_out	boolean	Denotes a Flexgroup.
share_count	integer	The number of shares in the application component.
storage_service	storage_service	
tiering	nas_application_components_tiering	application-components.tiering
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]}

app_cifs_access

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

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

app_nfs_access

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

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

protection_type

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

nas

A generic NAS application.

Name	Type	Description
application_components	array[application_components]	

Name	Type	Description
cifs_access	array[app_cifs_access]	The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.
nfs_access	array[app_nfs_access]	The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.
protection_type	protection_type	

performance

Name	Type	Description
storage_service	storage_service	

hosts

Name	Type	Description
nqn	string	The host NQN.

zapp_nvme_components_subsystem

components.subsystem

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

zapp_nvme_components_tiering

application-components.tiering

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

components

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

local

Name	Type	Description
name	string	The local RPO of the application.

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

rpo

Name	Type	Description
local	local	

zapp_nvme

An NVME application.

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

storage_service

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

archive_log

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

db

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

storage_service

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

ora_home

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

storage_service

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

redo_log

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

oracle_on_nfs

Oracle using NFS.

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

Name	Type	Description
ora_home	ora_home	
protection_type	protection_type	
redo_log	redo_log	

oracle_on_san_new_igroups

The list of initiator groups to create.

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

oracle_on_san

Oracle using SAN.

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

Name	Type	Description
protection_type	protection_type	
redo_log	redo_log	

storage_service

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

grid_binary

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

storage_service

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

oracle_crs

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

oracle_rac_on_nfs

Oracle RAC using NFS.

Name	Type	Description
archive_log	archive_log	
db	db	

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

db_sids

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

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

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

oracle_rac_on_san

Oracle RAC using SAN.

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

local

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

remote

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

rpo

Name	Type	Description
local	local	
remote	remote	

components

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

local

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

remote

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

rpo

Name	Type	Description
components	array[components]	

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

zapp_s3_bucket_application_components_access_policies_conditions

conditions

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

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

Name	Type	Description
actions	array[string]	
conditions	array[zapp_s3_bucket_application_components_access_policies_conditions]	conditions.
effect	string	Allow or Deny Access.
principal	array[string]	
resources	array[string]	
sid	string	Statement Identifier.

exclude_aggregates

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

zapp_s3_bucket_application_components

The list of application components to be created.

Name	Type	Description
access_policies	array[zapp_s3_bucket_application_components_access_policies]	The list of S3 objectstore policies to be created.
capacity_tier	boolean	Prefer lower latency storage under similar media costs.
comment	string	Object Store Server Bucket Description Usage: <(size 1..256)>
exclude_aggregates	array[exclude_aggregates]	
name	string	The name of the application component.
qos	qos	
size	integer	The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB MB GB TB PB]}
storage_service	storage_service	
uuid	string	Object Store Server Bucket UUID Usage: <UUID>

zapp_s3_bucket

A generic S3 bucket application.

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

san_application_components_tiering

application-components.tiering

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

application_components

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

san_new_igroups

The list of initiator groups to create.

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

san

A generic SAN application.

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

storage_service

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

db

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

storage_service

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

log

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

sql_on_san_new_igroups

The list of initiator groups to create.

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

storage_service

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

temp_db

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

sql_on_san

Microsoft SQL using SAN.

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

access

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

sql_on_smb

Microsoft SQL using SMB.

Name	Type	Description
access	access	
db	db	
log	log	
protection_type	protection_type	

Name	Type	Description
server_cores_count	integer	The number of server cores for the DB.
temp_db	temp_db	

iops

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

latency

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

snapshot

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

space

Name	Type	Description
available	integer	<p>The available amount of space left in the application component. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.</p> <ul style="list-style-type: none"> • readOnly: 1
logical_used	integer	<p>The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used.</p>
provisioned	integer	<p>The originally requested amount of space that was provisioned for the application component.</p>
reserved_unused	integer	<p>The amount of space reserved for system features such as Snapshot copies that has not yet been used.</p>
savings	integer	<p>The amount of space saved by all enabled space saving features.</p>
used	integer	<p>The amount of space that is currently being used by the application component. Note that this includes any space reserved by the system for features such as Snapshot copies.</p>
used_excluding_reserves	integer	<p>The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies.</p>
used_percent	integer	<p>The percentage of the originally provisioned space that is currently being used by the application component.</p>

storage_service

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

components

Name	Type	Description
iops	iops	
latency	latency	
name	string	Component Name.
shared_storage_pool	boolean	An application component is considered to use a shared storage pool if storage elements for other components reside on the same aggregate as storage elements for this component.
snapshot	snapshot	
space	space	
statistics_incomplete	boolean	If not all storage elements of the application component are currently available, the returned statistics might only include data from those elements that were available.
storage_service	storage_service	
uuid	string	Component UUID.

iops

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

Name	Type	Description
total	integer	The total number of IOPS being used by the application.

latency

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

space

Name	Type	Description
available	integer	<p>The available amount of space left in the application. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.</p> <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6
logical_used	integer	The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used.
provisioned	integer	The originally requested amount of space that was provisioned for the application.
reserved_unused	integer	The amount of space reserved for system features such as Snapshot copies that has not yet been used.

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

statistics

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

svm

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

self_link

template

Name	Type	Description
name	string	The name of the template that was used to provision this application.
protocol	string	The protocol access of the template that was used to provision this application.
version	integer	<p>The version of the template that was used to provision this application. The template version changes only if the layout of the application changes over time. For example, redo logs in Oracle RAC templates were updated and provisioned differently in DATA ONTAP 9.3.0 compared to prior releases, so the version number was increased. If layouts change in the future, the changes will be documented along with the corresponding version numbers.</p> <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6

storage_service

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

desktops

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

hyper_v_access

Name	Type	Description
service_account	string	Hyper-V service account.

vdi_on_nas

A VDI application using NAS.

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

vdi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	
name	string	The name of the new initiator group.
protocol	string	The protocol of the new initiator group.

vdi_on_san

A VDI application using SAN.

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

storage_service

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

datastore

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

vsi_on_nas

A VSI application using NAS.

Name	Type	Description
datastore	datastore	
hyper_v_access	hyper_v_access	

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

vsi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	
name	string	The name of the new initiator group.
protocol	string	The protocol of the new initiator group.

vsi_on_san

A VSI application using SAN.

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

application

Applications

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

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

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

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve application components

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

Introduced In: 9.6

Retrieves application components.

Overview

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

Query examples

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

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

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

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

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

Learn more

- [DOC /application](#)

Parameters

Name	Type	In	Required	Description
application.uuid	string	path	True	Application UUID
uuid	string	query	False	Filter by UUID
name	string	query	False	Filter by name
storage_service.name	string	query	False	Filter by storage_service.name
storage_service.uuid	string	query	False	Filter by storage_service.uuid
fields	array[string]	query	False	Specify the fields to return.

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

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

Example response

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

```

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

```



```

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

```

```

    }
  ],
  "protection_groups": [
    {
      "name": "string",
      "rpo": {
        "local": {
          "description": "string",
          "name": "string"
        },
        "remote": {
          "description": "string",
          "name": "string"
        }
      },
      "uuid": "string"
    }
  ],
  "san_access": [
    {
      "backing_storage": {
        "type": "string",
        "uuid": "string"
      },
      "lun_mappings": [
        {
          "fc": [
            {
              "interface": {
                "_links": {
                  "self": {
                    "href": "/api/resourceLink"
                  }
                }
              },
              "name": "lif1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
              "wwpn": "20:00:00:50:56:b4:13:a8"
            }
          ]
        },
        {
          "name": "string",
          "uuid": "string"
        }
      ]
    }
  ],
  "igroup": {
    "initiators": [
      "string"
    ],
    "name": "string",
    "uuid": "string"
  }
}

```

```

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

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

application

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

application_lun_object

LUN object

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

application_namespace_object

Namespace object

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

application_volume_object

Volume object

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

application_backing_storage

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

backing_storage

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

permissions

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

server

Name	Type	Description
name	string	Server name

share

Name	Type	Description
name	string	Share name

application_cifs_properties

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

export_policy

Name	Type	Description
name	string	Export policy name

permissions

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

application_nfs_properties

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

self

Name	Type	Description
self	href	

_links

Name	Type	Description
self	self	

hosts

Name	Type	Description
_links	_links	
nqn	string	Host

subsystem

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

subsystem_map

Subsystem map object

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

Name	Type	Description
subsystem	subsystem	

application_nvme_access

Application NVME access

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

local

Name	Type	Description
description	string	A detailed description of the local RPO. This includes details on the Snapshot copy schedule.
name	string	The local RPO of the component. This indicates how often component Snapshot copies are automatically created.

remote

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

rpo

Name	Type	Description
local	local	
remote	remote	

application_protection_groups

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

fc_interface_reference

An FC interface.

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

application_san_access_fcp_endpoint

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

Name	Type	Description
interface	fc_interface_reference	An FC interface.

igroup

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

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

application_san_access_iscsi_endpoint

An iSCSI access endpoint for the LUN.

Name	Type	Description
_links	_links	
interface	interface	
port	integer	The TCP port number of the iSCSI access endpoint.

application_lun_mapping_object

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

application_san_access

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

storage_service

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

svm

Name	Type	Description
name	string	SVM name
uuid	string	SVM UUID

application_component

Application component

Name	Type	Description
_links	_links	
application	application	
backing_storage	application_backing_storage	
cifs_access	array[application_cifs_properties]	
file_system	string	Defines the type of file system that will be installed on this application component.
host_management_url	string	Host management URL
host_name	string	L2 Host FQDN
name	string	Application component name
nfs_access	array[application_nfs_properties]	
nvme_access	array[application_nvme_access]	
protection_groups	array[application_protection_groups]	
san_access	array[application_san_access]	
storage_service	storage_service	
svm	svm	

Name	Type	Description
uuid	string	The application component UUID. Valid in URL.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve application component Snapshot copies

GET

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

Introduced In: 9.6

Retrieves Snapshot copies of an application component.

This endpoint is only supported for Maxdata template applications.

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

Learn more

- [DOC /application/applications/{application.uuid}/snapshots](#)
- [GET /application/applications/{uuid}/snapshots](#)
- [DOC /application](#)

Parameters

Name	Type	In	Required	Description
application.uuid	string	path	True	Application UUID
component.uuid	string	path	True	Application Component UUID
component.name	string	query	False	Filter by Application Component Name
uuid	string	query	False	Filter by uuid
name	string	query	False	Filter by name
consistency_type	string	query	False	Filter by consistency_type
comment	string	query	False	Filter by comment
create_time	string	query	False	Filter by create_time
is_partial	string	query	False	Filter by is_partial
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

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

Example response

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



```
]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

application

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

component

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

svm

Name	Type	Description
name	string	SVM Name
uuid	string	SVM UUID

application_component_snapshot

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Create an application component Snapshot copy

POST

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

Introduced In: 9.6

Creates a Snapshot copy of an application component.

This endpoint is only supported for Maxdata template applications.

Required properties

- name

Recommended optional properties

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

Learn more

- [DOC /application/applications/{application.uuid}/snapshots](#)
- [GET /application/applications/{uuid}/snapshots](#)
- [DOC /application](#)

Parameters

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

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

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

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

Example request

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

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

application

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

component

Name	Type	Description
name	string	Component Name
uuid	string	Component UUID

svm

Name	Type	Description
name	string	SVM Name
uuid	string	SVM UUID

application_component_snapshot

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

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

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete an application component Snapshot copy

DELETE

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

Introduced In: 9.6

Delete a Snapshot copy of an application component.

This endpoint is only supported for Maxdata template applications.

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

Learn more

- [DOC /application/applications/{application.uuid}/snapshots](#)
- [DELETE /application/applications/{application.uuid}/snapshots/{uuid}](#)
- [DOC /application](#)

Parameters

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

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve a Snapshot copy for a specific application component

GET

/application/applications/{application.uuid}/components/{component.uuid}/snapshot

s/{uuid}

Introduced In: 9.6

Retrieve a Snapshot copy of an application component.

This endpoint is only supported for Maxdata template applications.

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

Learn more

- [DOC /application/applications/{application.uuid}/snapshots](#)
- [GET /application/applications/{uuid}/snapshots](#)
- [DOC /application](#)

Parameters

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

Response

Status: 200, Ok

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

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

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

application

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

component

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

svm

Name	Type	Description
name	string	SVM Name
uuid	string	SVM UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Restore an application component Snapshot copy

POST

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

Introduced In: 9.6

Restore a Snapshot copy of an application component.

This endpoint is only supported for Maxdata template applications.

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

Learn more

- [DOC /application/applications/{application.uuid}/snapshots](#)
- [POST /application/applications/{application.uuid}/snapshots/{uuid}/restore](#)
- [DOC /application](#)
- [Asynchronous operations](#)

Parameters

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

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve an application component

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

Introduced In: 9.6

Retrieves an application component.

Overview

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

Access

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

Learn more

- [DOC /application](#)

Parameters

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

Response

Status: 200, Ok

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

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

Example response

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

```

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

```

```

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

```

```

    "remote": {
      "description": "string",
      "name": "string"
    }
  },
  "uuid": "string"
}
],
"san_access": [
  {
    "backing_storage": {
      "type": "string",
      "uuid": "string"
    },
    "lun_mappings": [
      {
        "fc": [
          {
            "interface": {
              "_links": {
                "self": {
                  "href": "/api/resourcelink"
                }
              },
            },
            "name": "lif1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
            "wwpn": "20:00:00:50:56:b4:13:a8"
          }
        ],
      },
    ],
    "igroup": {
      "initiators": [
        "string"
      ],
      "name": "string",
      "uuid": "string"
    },
    "iscsi": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
      },
      {
        "interface": {
          "_links": {

```

```

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

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

application

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

application_lun_object

LUN object

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

application_namespace_object

Namespace object

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

Name	Type	Description
size	integer	Namespace size
uuid	string	Namespace UUID

application_volume_object

Volume object

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

application_backing_storage

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

backing_storage

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

permissions

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

server

Name	Type	Description
name	string	Server name

share

Name	Type	Description
name	string	Share name

application_cifs_properties

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

export_policy

Name	Type	Description
name	string	Export policy name

permissions

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

application_nfs_properties

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

Name	Type	Description
permissions	array[permissions]	

self

Name	Type	Description
self	href	

_links

Name	Type	Description
self	self	

hosts

Name	Type	Description
_links	_links	
nqn	string	Host

subsystem

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

subsystem_map

Subsystem map object

Name	Type	Description
anagrpId	string	Subsystem ANA group ID
nsid	string	Subsystem namespace ID
subsystem	subsystem	

application_nvme_access

Application NVME access

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

local

Name	Type	Description
description	string	A detailed description of the local RPO. This includes details on the Snapshot copy schedule.
name	string	The local RPO of the component. This indicates how often component Snapshot copies are automatically created.

remote

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

rpo

Name	Type	Description
local	local	
remote	remote	

application_protection_groups

Name	Type	Description
name	string	Protection group name
rpo	rpo	

Name	Type	Description
uuid	string	Protection group UUID

fc_interface_reference

An FC interface.

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

application_san_access_fcp_endpoint

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

Name	Type	Description
interface	fc_interface_reference	An FC interface.

igroup

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

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

interface

Name	Type	Description
_links	_links	

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

application_san_access_iscsi_endpoint

An iSCSI access endpoint for the LUN.

Name	Type	Description
_links	_links	
interface	interface	
port	integer	The TCP port number of the iSCSI access endpoint.

application_lun_mapping_object

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

application_san_access

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

storage_service

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

svm

Name	Type	Description
name	string	SVM name
uuid	string	SVM UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Manage application Snapshot copies

Application applications application.uuid snapshots endpoint overview

Overview

Applications support Snapshot copies across all member storage elements. These Snapshot copies can be created and restored at any time or as scheduled. Most applications have hourly Snapshot copies enabled by default, unless the RPO setting is overridden during the creation of the application. An application Snapshot copy can be flagged as either *application consistent*, or *crash consistent*. From an ONTAP perspective, there is no difference between these two consistency types. These types are available for record keeping so that Snapshot copies taken after the application is quiesced (application consistent) can be tracked separately from

those Snapshot copies taken without first quiescing the application (crash consistent). By default, all application Snapshot copies are flagged to be *crash consistent*, and Snapshot copies taken at a scheduled time are also considered *crash consistent*.

The functionality provided by these APIs is not integrated with the host application. Snapshot copies have limited value without host coordination, so the use of the SnapCenter Backup Management suite is recommended to ensure correct interaction between host applications and ONTAP.

Retrieve an application Snapshot copy

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

Introduced In: 9.6

Retrieves Snapshot copies of an application.

Query examples

The following query returns all Snapshot copies from May 4, 2017 EST. For readability, the colon (:) is left in this example. For an actual call, they should be escaped as %3A.

```
GET
/application/applications/{application.uuid}/snapshots?create_time=2017-
05-04T00:00:00-05:00..2017-04T23:59:59-05:00
```

The following query returns all Snapshot copies that have been flagged as *application consistent*.

```
GET
/application/applications/{application.uuid}/snapshots?consistency_type=ap
plication
```

Learn more

- [DOC /application/applications/{application.uuid}/snapshots](#)
- [DOC /application](#)

Parameters

Name	Type	In	Required	Description
application.uuid	string	path	True	Application UUID
uuid	string	query	False	Filter by UUID
name	string	query	False	Filter by name
consistency_type	string	query	False	Filter by consistency_type

Name	Type	In	Required	Description
components.name	string	query	False	Filter by components.name
components.uuid	string	query	False	Filter by components.uuid
comment	string	query	False	Filter by comment
create_time	string	query	False	Filter by create_time
is_partial	string	query	False	Filter by is_partial
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[application_snapshot]	

Example response

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

```
    "uuid": "string"
  }
]
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

application

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

components

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

svm

Name	Type	Description
name	string	SVM name
uuid	string	SVM UUID

application_snapshot

Name	Type	Description
_links	_links	
application	application	
comment	string	Comment. Valid in POST.
components	array[components]	
consistency_type	string	Consistency type. This is for categorization purposes only. A Snapshot copy should not be set to 'application consistent' unless the host application is quiesced for the Snapshot copy. Valid in POST.
create_time	string	Creation time
is_partial	boolean	A partial Snapshot copy means that not all volumes in an application component were included in the Snapshot copy.
name	string	The Snapshot copy name. Valid in POST.
svm	svm	
uuid	string	The Snapshot copy UUID. Valid in URL.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Create an application Snapshot copy

POST /application/applications/{application.uuid}/snapshots

Introduced In: 9.6

Creates a Snapshot copy of the application.

Required properties

- name

Recommended optional properties

- consistency_type - Track whether this snapshot is *application* or *crash* consistent.

Learn more

- [DOC /application/applications/{application.uuid}/snapshots](#)
- [DOC /application](#)

Parameters

Name	Type	In	Required	Description
application.uuid	string	path	True	Application UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
application	application	
comment	string	Comment. Valid in POST.
components	array[components]	

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

Example request

```
{
  "application": {
    "name": "string",
    "uuid": "string"
  },
  "comment": "string",
  "components": [
    {
      "name": "string",
      "uuid": "string"
    }
  ],
  "consistency_type": "string",
  "create_time": "string",
  "name": "string",
  "svm": {
    "name": "string",
    "uuid": "string"
  },
  "uuid": "string"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

application

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

components

Name	Type	Description
name	string	Component name
uuid	string	Component UUID

svm

Name	Type	Description
name	string	SVM name
uuid	string	SVM UUID

application_snapshot

Name	Type	Description
application	application	
comment	string	Comment. Valid in POST.
components	array[components]	

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

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Delete an application Snapshot copy

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

Introduced In: 9.6

Delete a Snapshot copy of an application

Query examples

Individual Snapshot copies can be destroyed with no query parameters, or a range of Snapshot copies can be destroyed at one time using a query.

The following query deletes all application Snapshot copies created before May 4, 2017

```
DELETE
/application/applications/{application.uuid}/snapshots?create_time=<2017-
05-04T00:00:00-05:00
```

Parameters

Name	Type	In	Required	Description
application.uuid	string	path	True	Application UUID
uuid	string	path	True	Snapshot copy UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve an application Snapshot copy

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

Introduced In: 9.6

Retrieve a Snapshot copy of an application component.

This endpoint is only supported for Maxdata template applications.

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

Learn more

- [DOC /application/applications/{application.uuid}/snapshots](#)
- [GET /application/applications/{uuid}/snapshots](#)
- [DOC /application](#)

Parameters

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

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
application	application	
comment	string	Comment. Valid in POST.
components	array[components]	
consistency_type	string	Consistency type. This is for categorization purposes only. A Snapshot copy should not be set to 'application consistent' unless the host application is quiesced for the Snapshot copy. Valid in POST.
create_time	string	Creation time

Name	Type	Description
is_partial	boolean	A partial Snapshot copy means that not all volumes in an application component were included in the Snapshot copy.
name	string	The Snapshot copy name. Valid in POST.
svm	svm	
uuid	string	The Snapshot copy UUID. Valid in URL.

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

application

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

components

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

svm

Name	Type	Description
name	string	SVM name
uuid	string	SVM UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Restore an application Snapshot copy

POST /application/applications/{application.uuid}/snapshots/{uuid}/restore

Introduced In: 9.6

Restore an application snapshot

Restoring an application Snapshot copy reverts all storage elements in the Snapshot copy to the state in which the Snapshot copy was in when the Snapshot copy was taken. This restoration does not apply to access settings that might have changed since the Snapshot copy was created.

Learn more

- [DOC /application](#)
- [Asynchronous operations](#)

Parameters

Name	Type	In	Required	Description
application.uuid	string	path	True	Application UUID
uuid	string	path	True	Snapshot copy UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete an application and all associated data

DELETE /application/applications/{uuid}

Introduced In: 9.6

Deletes an application and all associated data.

Warning - this deletes it all, including your data

This deletes everything created with the application, including any volumes, LUNs, NFS export policies, CIFS shares, and initiator groups. Initiator groups are only destroyed if they were created as part of an application and are no longer in use by other applications.

Learn more

- [DOC /application](#)
- [Asynchronous operations](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Application UUID
delete_data	boolean	query	False	<p>By default, deleting an application deletes all of the application's data. By setting this parameter to "false", the application's data is preserved, but can no longer be managed through application APIs.</p> <ul style="list-style-type: none">• Introduced in: 9.8• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve an application

GET /application/applications/{uuid}

Introduced In: 9.6

Retrieves an application

Expensive properties

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

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

Property overview

An application includes three main groups or properties.

- Generic properties - such as the `name`, `template.name`, and `state` of the application. These properties are all inexpensive to retrieve and their meaning is consistent for every type of application.
- `statistics.*` - application statistics report live usage data about the application and its components. Various space and IOPS details are included at both the application level and at a per component level. The application model includes a detailed description of each property. These properties are slightly more expensive than the generic properties because live data must be collected from every storage element in the application.
- `<template>` - the property corresponding to the value of the `template.name` returns the contents of the application in the same layout that was used to provision the application. This information is very expensive to retrieve because it requires collecting information about all the storage and access settings for every element of the application. There are a few notable limitations to what can be returned in the `<template>` section:
 - The `new_igroups` array of many SAN templates is not returned by GET. This property allows igroup creation in the same call that creates an application, but is not a property of the application itself. The `new_igroups` array is allowed during PATCH operations, but that does not modify the `new_igroups` of the application. It is another way to allow igroup creation while updating the application to use a different igroup.
 - The `vdi_on_san` and `vdi_on_nas` `desktops.count` property is rounded to the nearest 1000 during creation, and is reported with that rounding applied.
 - The `mongo_db_on_san` `dataset.element_count` property is rounded up to an even number, and is reported with that rounding applied.
 - The `sql_on_san` and `sql_on_smb` `server_cores_count` property is limited to 8 for GET operations. Higher values are accepted by POST, but the impact of the `server_cores_count` property on the application layout currently reaches its limit at 8.

Learn more

- [DOC /application](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Application UUID

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

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

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

Name	Type	Description
vsi_on_nas	vsi_on_nas	A VSI application using NAS.
vsi_on_san	vsi_on_san	A VSI application using SAN.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    },
    "snapshots": {
      "href": "/api/resourcelink"
    }
  },
  "cg_on_san": {
    "cgs": [
      {
        "lun_containers": [
          {
            "luns": [
              {
                "igroups": [
                  {
                    "name": "string"
                  }
                ],
                "name": "string",
                "os_type": "string"
              }
            ],
            "name": "string"
          }
        ],
        "name": "string",
        "protection_type": {
          "local_policy": "string",
          "remote_rpo": "string"
        },
        "qos": {
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "storage_service": {
          "name": "string"
        },
        "tiering": {
          "policy": "string"
        }
      }
    ]
  }
}
```

```

    }
  }
]
},
"creation_timestamp": "string",
"delete_data": null,
"generation": 0,
"maxdata_on_san": {
  "app_type": "string",
  "application_components": [
    {}
  ],
  "metadata": [
    {
      "key": "string",
      "value": "string"
    }
  ],
  "ocsm_url": "string",
  "os_type": "string"
},
"name": "string",
"nas": {
  "application_components": [
    {}
  ],
  "cifs_access": [
    {
      "access": "string",
      "user_or_group": "string"
    }
  ],
  "nfs_access": [
    {
      "access": "string",
      "host": "string"
    }
  ],
  "protection_type": {
    "local_policy": "string",
    "local_rpo": "string",
    "remote_rpo": "string"
  }
},
"nvme": {
  "components": [

```

```

    {
      "name": "string",
      "os_type": "string",
      "performance": {
        "storage_service": {
          "name": "string"
        }
      },
      "qos": {
        "policy": {
          "name": "string",
          "uuid": "string"
        }
      },
      "subsystem": {
        "hosts": [
          {
            "nqn": "string"
          }
        ],
        "name": "string",
        "os_type": "string",
        "uuid": "string"
      },
      "tiering": {
        "policy": "string"
      }
    }
  ],
  "os_type": "string",
  "rpo": {
    "local": {
      "name": "string",
      "policy": "string"
    }
  }
},
"protection_granularity": "string",
"rpo": {
  "components": [
    {
      "name": "string",
      "rpo": {
        "local": {
          "description": "string",
          "name": "string"
        }
      }
    }
  ]
}

```

```

        },
        "remote": {
            "description": "string",
            "name": "string"
        }
    },
    "uuid": "string"
}
],
"local": {
    "description": "string",
    "name": "string"
},
"remote": {
    "description": "string",
    "name": "string"
}
},
"s3_bucket": {
    "application_components": [
        {}
    ]
},
"san": {
    "application_components": [
        {
            "igroup_name": "string",
            "name": "string",
            "os_type": "string",
            "qos": {
                "policy": {
                    "name": "string",
                    "uuid": "string"
                }
            },
            "storage_service": {
                "name": "string"
            },
            "tiering": {
                "policy": "string"
            }
        }
    ]
},
"os_type": "string",
"protection_type": {
    "local_policy": "string",

```

```

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

```

```

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

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	
snapshots	href	

igroups

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

zapp_cg_on_san_cgs_lun_containers_luns

The list of LUNs to be created.

Name	Type	Description
count	integer	The number of LUNs.
igroups	array[igroups]	
name	string	The name/prefix of LUNs.
os_type	string	The name of the host OS.

zapp_cg_on_san_cgs_lun_containers

LUN containers.

Name	Type	Description
luns	array[zapp_cg_on_san_cgs_lun_containers_luns]	The list of LUNs to be created.

Name	Type	Description
name	string	The name or prefix of container (Volume). If the container exists, it becomes a part of the CG. If not, a new container is created.

protection_type

Name	Type	Description
local_policy	string	The Snapshot copy policy to apply to a CG.
remote_rpo	string	The remote RPO of the CG.

policy

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

qos

Name	Type	Description
policy	policy	

storage_service

Name	Type	Description
name	string	The storage service of the CG.

object_stores

zapp_cg_on_san_cgs_tiering

tiering

Name	Type	Description
policy	string	The storage tiering type of the CG.

zapp_cg_on_san_cgs

The list of CGs to be created.

Name	Type	Description
lun_containers	array[zapp_cg_on_san_cgs_lun_containers]	LUN containers.
name	string	The name of the consistency group.
protection_type	protection_type	
qos	qos	
storage_service	storage_service	
tiering	zapp_cg_on_san_cgs_tiering	tiering

zapp_cg_on_san

A consistency group (CG) application using SAN.

Name	Type	Description
cgs	array[zapp_cg_on_san_cgs]	The list of CGs to be created.

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application component.
value	string	Value associated with the key.

protection_type

Name	Type	Description
local_rpo	string	The local rpo of the application component.
remote_rpo	string	The remote rpo of the application component.

storage_service

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

maxdata_on_san_application_components_tiering

tiering

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

maxdata_on_san_application_components

The list of application components to be created.

Name	Type	Description
file_system	string	Defines the type of file system that will be installed on this application component.
host_management_url	string	The host management URL for this application component.
host_name	string	FQDN of the L2 host that contains the hot tier of this application component.
igroup_name	string	The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup.
lun_count	integer	The number of LUNs in the application component.
metadata	array[metadata]	
name	string	The name of the application component.
protection_type	protection_type	

Name	Type	Description
storage_service	storage_service	
tiering	maxdata_on_san_application_components_tiering	tiering
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]}

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application.
value	string	Value associated with the key.

maxdata_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

maxdata_on_san

MAX Data application using SAN.

Name	Type	Description
app_type	string	Type of the application that is being deployed on the L2.
application_components	array[maxdata_on_san_application_components]	The list of application components to be created.
metadata	array[metadata]	
ocsm_url	string	The OnCommand System Manager URL for this application.
os_type	string	The name of the host OS running the application.

storage_service

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

dataset

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

mongo_db_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

protection_type

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

secondary_igroups

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

mongo_db_on_san

MongoDB using SAN.

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

export_policy

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

component

Name	Type	Description
name	string	Name of the source component.

svm

Name	Type	Description
name	string	Name of the source SVM.

origin

Name	Type	Description
component	component	
svm	svm	

flexcache

Name	Type	Description
origin	origin	

nas_application_components_tiering

application-components.tiering

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

application_components

Name	Type	Description
export_policy	export_policy	
flexcache	flexcache	
name	string	The name of the application component.
qos	qos	
scale_out	boolean	Denotes a Flexgroup.
share_count	integer	The number of shares in the application component.
storage_service	storage_service	
tiering	nas_application_components_tiering	application-components.tiering
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]}

app_cifs_access

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

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

app_nfs_access

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

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

protection_type

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

nas

A generic NAS application.

Name	Type	Description
application_components	array[application_components]	
cifs_access	array[app_cifs_access]	The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.
nfs_access	array[app_nfs_access]	The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.
protection_type	protection_type	

performance

Name	Type	Description
storage_service	storage_service	

hosts

Name	Type	Description
nqn	string	The host NQN.

zapp_nvme_components_subsystem

components.subsystem

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

zapp_nvme_components_tiering

application-components.tiering

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

components

Name	Type	Description
name	string	The name of the application component.
namespace_count	integer	The number of namespaces in the component.

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

local

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

rpo

Name	Type	Description
local	local	

zapp_nvme

An NVME application.

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

storage_service

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

archive_log

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

db

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

storage_service

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

ora_home

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

storage_service

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

redo_log

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

oracle_on_nfs

Oracle using NFS.

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

oracle_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

oracle_on_san

Oracle using SAN.

Name	Type	Description
archive_log	archive_log	
db	db	

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

storage_service

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

grid_binary

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

storage_service

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

oracle_crs

Name	Type	Description
copies	integer	The number of CRS volumes.

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

oracle_rac_on_nfs

Oracle RAC using NFS.

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

db_sids

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

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

oracle_rac_on_san

Oracle RAC using SAN.

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

local

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

remote

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

rpo

Name	Type	Description
local	local	
remote	remote	

components

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

local

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

remote

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

rpo

Name	Type	Description
components	array[components]	

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

zapp_s3_bucket_application_components_access_policies_conditions

conditions

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

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

Name	Type	Description
actions	array[string]	
conditions	array[zapp_s3_bucket_application_components_access_policies_conditions]	conditions.
effect	string	Allow or Deny Access.
principal	array[string]	
resources	array[string]	
sid	string	Statement Identifier.

exclude_aggregates

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

zapp_s3_bucket_application_components

The list of application components to be created.

Name	Type	Description
access_policies	array[zapp_s3_bucket_application_components_access_policies]	The list of S3 objectstore policies to be created.
capacity_tier	boolean	Prefer lower latency storage under similar media costs.
comment	string	Object Store Server Bucket Description Usage: <(size 1..256)>
exclude_aggregates	array[exclude_aggregates]	
name	string	The name of the application component.
qos	qos	
size	integer	The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB MB GB TB PB]}
storage_service	storage_service	
uuid	string	Object Store Server Bucket UUID Usage: <UUID>

zapp_s3_bucket

A generic S3 bucket application.

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

san_application_components_tiering

application-components.tiering

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

application_components

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

san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

san

A generic SAN application.

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

storage_service

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

db

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

storage_service

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

log

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

sql_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

storage_service

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

temp_db

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

sql_on_san

Microsoft SQL using SAN.

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

access

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

sql_on_smb

Microsoft SQL using SMB.

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

iops

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

latency

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

snapshot

Name	Type	Description
reserve	integer	The amount of space reserved by the system for Snapshot copies.

Name	Type	Description
used	integer	The amount of spacing currently in use by the system to store Snapshot copies.

space

Name	Type	Description
available	integer	<p>The available amount of space left in the application component. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.</p> <ul style="list-style-type: none"> • readOnly: 1
logical_used	integer	The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used.
provisioned	integer	The originally requested amount of space that was provisioned for the application component.
reserved_unused	integer	The amount of space reserved for system features such as Snapshot copies that has not yet been used.
savings	integer	The amount of space saved by all enabled space saving features.
used	integer	The amount of space that is currently being used by the application component. Note that this includes any space reserved by the system for features such as Snapshot copies.

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

storage_service

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

components

Name	Type	Description
iops	iops	
latency	latency	
name	string	Component Name.
shared_storage_pool	boolean	An application component is considered to use a shared storage pool if storage elements for other components reside on the same aggregate as storage elements for this component.
snapshot	snapshot	
space	space	
statistics_incomplete	boolean	If not all storage elements of the application component are currently available, the returned statistics might only include data from those elements that were available.

Name	Type	Description
storage_service	storage_service	
uuid	string	Component UUID.

iops

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

latency

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

space

Name	Type	Description
available	integer	<p>The available amount of space left in the application. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.</p> <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6

Name	Type	Description
logical_used	integer	The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used.
provisioned	integer	The originally requested amount of space that was provisioned for the application.
reserved_unused	integer	The amount of space reserved for system features such as Snapshot copies that has not yet been used.
savings	integer	The amount of space saved by all enabled space saving features.
used	integer	The amount of space that is currently being used by the application. Note that this includes any space reserved by the system for features such as Snapshot copies.
used_excluding_reserves	integer	The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies.
used_percent	integer	The percentage of the originally provisioned space that is currently being used by the application.

statistics

Name	Type	Description
components	array[components]	
iops	iops	
latency	latency	

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

svm

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

self_link

Name	Type	Description
self	href	

template

Name	Type	Description
_links	self_link	
name	string	The name of the template that was used to provision this application.
protocol	string	The protocol access of the template that was used to provision this application.

Name	Type	Description
version	integer	<p>The version of the template that was used to provision this application. The template version changes only if the layout of the application changes over time. For example, redo logs in Oracle RAC templates were updated and provisioned differently in DATA ONTAP 9.3.0 compared to prior releases, so the version number was increased. If layouts change in the future, the changes will be documented along with the corresponding version numbers.</p> <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6

storage_service

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

desktops

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

hyper_v_access

Name	Type	Description
service_account	string	Hyper-V service account.

vdi_on_nas

A VDI application using NAS.

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

vdi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

vdi_on_san

A VDI application using SAN.

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

storage_service

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

datastore

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

vsi_on_nas

A VSI application using NAS.

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

vsi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

vsi_on_san

A VSI application using SAN.

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

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Update application properties

PATCH /application/applications/{uuid}

Introduced In: 9.6

Updates the properties of an application.

Overview

Similar to creating an application, modification is done using the template properties of an application. The `storage_service`, `size`, and `igroup_name` of an application may be modified.

storage_service

Storage service modifications are processed in place, meaning that the storage can not be moved to a location with more performance headroom to accommodate the request. If the current backing storage of the application is in a location that can support increased performance, the QoS policies associated with the application will be modified to allow it. If not, an error will be returned. A storage service modification to a lower tier of performance is always allowed, but the reverse modification may not be supported if the cluster is over provisioned and the cluster is unlikely to be able to fulfil the original storage service.

size

Size modifications are processed in a variety of ways depending on the type of application. For NAS applications, volumes are grown or new volumes are added. For SAN applications, LUNs are grown, new LUNs are added to existing volumes, or new LUNs are added to new volumes. If new storage elements are created, they can be found using the [GET /application/applications/{application.uuid}/components](#) interface. The creation time of each storage object is included, and the newly created objects will use the same naming scheme as the previous objects. Resize follows the best practices associated with the type of application being expanded. Reducing the size of an application is not supported.

igroup_name

Modification of the igroup name allows an entire application to be mapped from one initiator group to another. Data access will be interrupted as the LUNs are unmapped from the original igroup and remapped to the new one.

Application state

During a modification, the `state` property of the application updates to indicate `modifying`. In `modifying` state, statistics are not available and Snapshot copy operations are not allowed. If the modification fails, it is possible for the application to be left in an inconsistent state, with the underlying ONTAP storage elements not matching across a component. When this occurs, the application is left in the `modifying` state until the command is either retried and succeeds or a call to restore the original state is successful.

Examples

1. Change the storage service of the database of the Oracle application to *extreme* and resize the redo logs to *100GB*.

```
{
  "oracle_on_nfs": {
    "db": {
      "storage_service": {
        "name": "extreme"
      }
    },
    "redo_log": {
      "size": "100GB"
    }
  }
}
```

2. Change the storage service, size, and igroup of a generic application by component name.

```
{
  "san": {
    "application_components": [
      {
        "name": "component1",
        "storage_service": {
          "name": "value"
        }
      },
      {
        "name": "component2",
        "size": "200GB"
      },
      {
        "name": "component3",
        "igroup_name": "igroup5"
      }
    ]
  }
}
```

Learn more

- [DOC /application](#)
- [Asynchronous operations](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Application UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
_links	_links	
cg_on_san	zapp_cg_on_san	A consistency group (CG) application using SAN.
creation_timestamp	string	The time when the application was created.

Name	Type	Description
generation	integer	The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented.
maxdata_on_san	maxdata_on_san	MAX Data application using SAN.
mongo_db_on_san	mongo_db_on_san	MongoDB using SAN.
nas	nas	A generic NAS application.
nvme	zapp_nvme	An NVME application.
oracle_on_nfs	oracle_on_nfs	Oracle using NFS.
oracle_on_san	oracle_on_san	Oracle using SAN.
oracle_rac_on_nfs	oracle_rac_on_nfs	Oracle RAC using NFS.
oracle_rac_on_san	oracle_rac_on_san	Oracle RAC using SAN.
protection_granularity	string	Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components.
rpo	rpo	
s3_bucket	zapp_s3_bucket	A generic S3 bucket application.
san	san	A generic SAN application.
sql_on_san	sql_on_san	Microsoft SQL using SAN.

Name	Type	Description
sql_on_smb	sql_on_smb	Microsoft SQL using SMB.
state	string	The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported.
statistics	statistics	
template	template	
uuid	string	Application UUID. This field is generated when the application is created.
vdi_on_nas	vdi_on_nas	A VDI application using NAS.
vdi_on_san	vdi_on_san	A VDI application using SAN.
vsi_on_nas	vsi_on_nas	A VSI application using NAS.
vsi_on_san	vsi_on_san	A VSI application using SAN.

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

igroups

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

zapp_cg_on_san_cgs_lun_containers_luns

The list of LUNs to be created.

Name	Type	Description
igroups	array[igroups]	

zapp_cg_on_san_cgs_lun_containers

LUN containers.

Name	Type	Description
luns	array[zapp_cg_on_san_cgs_lun_containers_luns]	The list of LUNs to be created.

protection_type

policy

qos

storage_service

Name	Type	Description
name	string	The storage service of the CG.

object_stores

zapp_cg_on_san_cgs_tiering

tiering

zapp_cg_on_san_cgs

The list of CGs to be created.

Name	Type	Description
lun_containers	array[zapp_cg_on_san_cgs_lun_containers]	LUN containers.
name	string	The name of the consistency group.
storage_service	storage_service	
tiering	zapp_cg_on_san_cgs_tiering	tiering

zapp_cg_on_san

A consistency group (CG) application using SAN.

Name	Type	Description
cgs	array[zapp_cg_on_san_cgs]	The list of CGs to be created.

metadata

protection_type

Name	Type	Description
local_rpo	string	The local rpo of the application component.
remote_rpo	string	The remote rpo of the application component.

storage_service

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

maxdata_on_san_application_components_tiering

tiering

maxdata_on_san_application_components

The list of application components to be created.

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

maxdata_on_san_new_igroups

The list of initiator groups to create.

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

Name	Type	Description
protocol	string	The protocol of the new initiator group.

maxdata_on_san

MAX Data application using SAN.

Name	Type	Description
application_components	array[maxdata_on_san_application_components]	The list of application components to be created.
new_igroups	array[maxdata_on_san_new_igroups]	The list of initiator groups to create.
ocsm_url	string	The OnCommand System Manager URL for this application.

storage_service

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

dataset

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

mongo_db_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	
name	string	The name of the new initiator group.

Name	Type	Description
os_type	string	The name of the host OS accessing the application. The default value is the host OS that is running the application.
protocol	string	The protocol of the new initiator group.

protection_type

Name	Type	Description
local_rpo	string	The local RPO of the application.

secondary_igroups

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

mongo_db_on_san

MongoDB using SAN.

Name	Type	Description
dataset	dataset	
new_igroups	array[mongo_db_on_san_new_igroups]	The list of initiator groups to create.
primary_igroup_name	string	The initiator group for the primary.
protection_type	protection_type	
secondary_igroups	array[secondary_igroups]	

export_policy

component

svm

origin

flexcache

nas_application_components_tiering

application-components.tiering

application_components

Name	Type	Description
name	string	The name of the application component.
storage_service	storage_service	
tiering	nas_application_components_tiering	application-components.tiering
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]}

app_cifs_access

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

app_nfs_access

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

nas

A generic NAS application.

Name	Type	Description
application_components	array[application_components]	
protection_type	protection_type	

performance

Name	Type	Description
storage_service	storage_service	

hosts

zapp_nvme_components_subsystem

components.subsystem

zapp_nvme_components_tiering

application-components.tiering

components

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

local

Name	Type	Description
name	string	The local RPO of the application.

rpo

Name	Type	Description
local	local	

zapp_nvme

An NVME application.

Name	Type	Description
components	array[components]	
rpo	rpo	

storage_service

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

archive_log

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

db

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

storage_service

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

ora_home

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

storage_service

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

redo_log

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

oracle_on_nfs

Oracle using NFS.

Name	Type	Description
archive_log	archive_log	
db	db	
ora_home	ora_home	
protection_type	protection_type	
redo_log	redo_log	

oracle_on_san_new_igroups

The list of initiator groups to create.

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

oracle_on_san

Oracle using SAN.

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

Name	Type	Description
new_igroups	array[oracle_on_san_new_igroups]	The list of initiator groups to create.
ora_home	ora_home	
protection_type	protection_type	
redo_log	redo_log	

storage_service

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

grid_binary

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

storage_service

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

oracle_crs

Name	Type	Description
storage_service	storage_service	

oracle_rac_on_nfs

Oracle RAC using NFS.

Name	Type	Description
archive_log	archive_log	
db	db	
grid_binary	grid_binary	
ora_home	ora_home	

Name	Type	Description
oracle_crs	oracle_crs	
protection_type	protection_type	
redo_log	redo_log	

db_sids

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

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

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

oracle_rac_on_san

Oracle RAC using SAN.

Name	Type	Description
archive_log	archive_log	
db	db	
db_sids	array[db_sids]	
grid_binary	grid_binary	

Name	Type	Description
new_igroups	array[oracle_rac_on_san_new_igroups]	The list of initiator groups to create.
ora_home	ora_home	
oracle_crs	oracle_crs	
protection_type	protection_type	
redo_log	redo_log	

local

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

remote

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

rpo

Name	Type	Description
local	local	
remote	remote	

components

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

local

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

remote

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

rpo

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

zapp_s3_bucket_application_components_access_policies_conditions

conditions

Name	Type	Description
delimiters	array[string]	
max_keys	array[integer]	
prefixes	array[string]	
source_ips	array[string]	
usernames	array[string]	

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

Name	Type	Description
actions	array[string]	
principal	array[string]	
resources	array[string]	

exclude_aggregates

zapp_s3_bucket_application_components

The list of application components to be created.

Name	Type	Description
name	string	The name of the application component.
size	integer	The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB MB GB TB PB]}
storage_service	storage_service	
uuid	string	Object Store Server Bucket UUID Usage: <UUID>

zapp_s3_bucket

A generic S3 bucket application.

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

san_application_components_tiering

application-components.tiering

application_components

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

san_new_igroups

The list of initiator groups to create.

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

Name	Type	Description
protocol	string	The protocol of the new initiator group.

san

A generic SAN application.

Name	Type	Description
application_components	array[application_components]	
new_igroups	array[san_new_igroups]	The list of initiator groups to create.
protection_type	protection_type	

storage_service

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

db

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

storage_service

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

log

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

sql_on_san_new_igroups

The list of initiator groups to create.

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

storage_service

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

temp_db

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

sql_on_san

Microsoft SQL using SAN.

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

Name	Type	Description
new_igroups	array[sql_on_san_new_igroups]	The list of initiator groups to create.
protection_type	protection_type	
temp_db	temp_db	

access

sql_on_smb

Microsoft SQL using SMB.

Name	Type	Description
db	db	
log	log	
protection_type	protection_type	
temp_db	temp_db	

iops

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

latency

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

snapshot

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

space

Name	Type	Description
available	integer	<p>The available amount of space left in the application component. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.</p> <ul style="list-style-type: none"> • readOnly: 1
logical_used	integer	The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used.
provisioned	integer	The originally requested amount of space that was provisioned for the application component.
reserved_unused	integer	The amount of space reserved for system features such as Snapshot copies that has not yet been used.
savings	integer	The amount of space saved by all enabled space saving features.

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

storage_service

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

components

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

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

iops

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

latency

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

space

Name	Type	Description
available	integer	<p>The available amount of space left in the application. Note that this field has limited meaning for SAN applications. Space may be considered used from ONTAP's perspective while the host filesystem still considers it available.</p> <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6

Name	Type	Description
logical_used	integer	The amount of space that would currently be used if no space saving features were enabled. For example, if compression were the only space saving feature enabled, this field would represent the uncompressed amount of space used.
provisioned	integer	The originally requested amount of space that was provisioned for the application.
reserved_unused	integer	The amount of space reserved for system features such as Snapshot copies that has not yet been used.
savings	integer	The amount of space saved by all enabled space saving features.
used	integer	The amount of space that is currently being used by the application. Note that this includes any space reserved by the system for features such as Snapshot copies.
used_excluding_reserves	integer	The amount of space that is currently being used, excluding any space that is reserved by the system for features such as Snapshot copies.
used_percent	integer	The percentage of the originally provisioned space that is currently being used by the application.

statistics

Name	Type	Description
components	array[components]	
iops	iops	
latency	latency	

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

self_link

template

Name	Type	Description
protocol	string	The protocol access of the template that was used to provision this application.
version	integer	<p>The version of the template that was used to provision this application. The template version changes only if the layout of the application changes over time. For example, redo logs in Oracle RAC templates were updated and provisioned differently in DATA ONTAP 9.3.0 compared to prior releases, so the version number was increased. If layouts change in the future, the changes will be documented along with the corresponding version numbers.</p> <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6

storage_service

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

desktops

Name	Type	Description
count	integer	The number of desktops to support.
storage_service	storage_service	

hyper_v_access

vdi_on_nas

A VDI application using NAS.

Name	Type	Description
desktops	desktops	
protection_type	protection_type	

vdi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	
name	string	The name of the new initiator group.
protocol	string	The protocol of the new initiator group.

vdi_on_san

A VDI application using SAN.

Name	Type	Description
desktops	desktops	

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

storage_service

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

datastore

Name	Type	Description
count	integer	The number of datastores to support.
storage_service	storage_service	

vsi_on_nas

A VSI application using NAS.

Name	Type	Description
datastore	datastore	
protection_type	protection_type	

vsi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

Name	Type	Description
name	string	The name of the new initiator group.
protocol	string	The protocol of the new initiator group.

vsi_on_san

A VSI application using SAN.

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

application

Applications

Name	Type	Description
_links	_links	
cg_on_san	zapp_cg_on_san	A consistency group (CG) application using SAN.
creation_timestamp	string	The time when the application was created.

Name	Type	Description
generation	integer	The generation number of the application. This indicates which features are supported on the application. For example, generation 1 applications do not support Snapshot copies. Support for Snapshot copies was added at generation 2. Any future generation numbers and their feature set will be documented.
maxdata_on_san	maxdata_on_san	MAX Data application using SAN.
mongo_db_on_san	mongo_db_on_san	MongoDB using SAN.
nas	nas	A generic NAS application.
nvme	zapp_nvme	An NVME application.
oracle_on_nfs	oracle_on_nfs	Oracle using NFS.
oracle_on_san	oracle_on_san	Oracle using SAN.
oracle_rac_on_nfs	oracle_rac_on_nfs	Oracle RAC using NFS.
oracle_rac_on_san	oracle_rac_on_san	Oracle RAC using SAN.
protection_granularity	string	Protection granularity determines the scope of Snapshot copy operations for the application. Possible values are "application" and "component". If the value is "application", Snapshot copy operations are performed on the entire application. If the value is "component", Snapshot copy operations are performed separately on the application components.
rpo	rpo	
s3_bucket	zapp_s3_bucket	A generic S3 bucket application.
san	san	A generic SAN application.
sql_on_san	sql_on_san	Microsoft SQL using SAN.

Name	Type	Description
sql_on_smb	sql_on_smb	Microsoft SQL using SMB.
state	string	The state of the application. For full functionality, applications must be in the online state. Other states indicate that the application is in a transient state and not all operations are supported.
statistics	statistics	
template	template	
uuid	string	Application UUID. This field is generated when the application is created.
vdi_on_nas	vdi_on_nas	A VDI application using NAS.
vdi_on_san	vdi_on_san	A VDI application using SAN.
vsi_on_nas	vsi_on_nas	A VSI application using NAS.
vsi_on_san	vsi_on_san	A VSI application using SAN.

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve application templates

GET /application/templates

Introduced In: 9.6

Retrieves application templates.

Query examples

The most useful queries on this API allows searches by name or protocol access. The following query returns all templates that are used to provision an Oracle application.

```
GET /application/templates?name=ora*
```

Similarly, the following query returns all templates that support SAN access.

```
GET /application/templates?protocol=san
```

Learn more

- [DOC /application](#)

Parameters

Name	Type	In	Required	Description
name	string	query	False	Filter by name
protocol	string	query	False	Filter by protocol
description	string	query	False	Filter by description
missing_prerequisites	string	query	False	Filter by missing_prerequisites

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[application_template]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cg_on_san": {
        "cgs": [
          {
            "lun_containers": [
              {
                "luns": [
                  {
                    "igroups": [
                      {
                        "name": "string"
                      }
                    ],
                    "name": "string",
                    "os_type": "string"
                  }
                ],
                "name": "string"
              }
            ],
            "name": "string",
            "protection_type": {
              "local_policy": "string",
              "remote_rpo": "string"
            },
            "qos": {
              "policy": {
                "name": "string",
                "uuid": "string"
              }
            }
          }
        ]
      }
    }
  ]
}
```

```

        }
    },
    "storage_service": {
        "name": "string"
    },
    "tiering": {
        "policy": "string"
    }
}
]
},
"description": "string",
"maxdata_on_san": {
    "app_type": "string",
    "application_components": [
        {}
    ],
    "metadata": [
        {
            "key": "string",
            "value": "string"
        }
    ],
    "ocsm_url": "string",
    "os_type": "string"
},
"missing_prerequisites": "string",
"name": "string",
"nas": {
    "application_components": [
        {}
    ],
    "cifs_access": [
        {
            "access": "string",
            "user_or_group": "string"
        }
    ],
    "nfs_access": [
        {
            "access": "string",
            "host": "string"
        }
    ],
    "protection_type": {
        "local_policy": "string",

```

```

        "local_rpo": "string",
        "remote_rpo": "string"
    }
},
"nvme": {
    "components": [
        {
            "name": "string",
            "os_type": "string",
            "performance": {
                "storage_service": {
                    "name": "string"
                }
            },
            "qos": {
                "policy": {
                    "name": "string",
                    "uuid": "string"
                }
            },
            "subsystem": {
                "hosts": [
                    {
                        "nqn": "string"
                    }
                ],
                "name": "string",
                "os_type": "string",
                "uuid": "string"
            },
            "tiering": {
                "policy": "string"
            }
        }
    ],
    "os_type": "string",
    "rpo": {
        "local": {
            "name": "string",
            "policy": "string"
        }
    },
    "protocol": "string",
    "s3_bucket": {
        "application_components": [

```

```

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

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

self_link

Name	Type	Description
self	href	

igroups

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

zapp_cg_on_san_cgs_lun_containers_luns

The list of LUNs to be created.

Name	Type	Description
count	integer	The number of LUNs.
igroups	array[igroups]	
name	string	The name/prefix of LUNs.
os_type	string	The name of the host OS.

zapp_cg_on_san_cgs_lun_containers

LUN containers.

Name	Type	Description
luns	array[zapp_cg_on_san_cgs_lun_containers_luns]	The list of LUNs to be created.
name	string	The name or prefix of container (Volume). If the container exists, it becomes a part of the CG. If not, a new container is created.

protection_type

Name	Type	Description
local_policy	string	The Snapshot copy policy to apply to a CG.
remote_rpo	string	The remote RPO of the CG.

policy

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

qos

Name	Type	Description
policy	policy	

storage_service

Name	Type	Description
name	string	The storage service of the CG.

object_stores

zapp_cg_on_san_cgs_tiering

tiering

Name	Type	Description
policy	string	The storage tiering type of the CG.

zapp_cg_on_san_cgs

The list of CGs to be created.

Name	Type	Description
lun_containers	array[zapp_cg_on_san_cgs_lun_containers]	LUN containers.
name	string	The name of the consistency group.
protection_type	protection_type	
qos	qos	
storage_service	storage_service	
tiering	zapp_cg_on_san_cgs_tiering	tiering

zapp_cg_on_san

A consistency group (CG) application using SAN.

Name	Type	Description
cgs	array[zapp_cg_on_san_cgs]	The list of CGs to be created.

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application component.
value	string	Value associated with the key.

protection_type

Name	Type	Description
local_rpo	string	The local rpo of the application component.

Name	Type	Description
remote_rpo	string	The remote rpo of the application component.

storage_service

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

maxdata_on_san_application_components_tiering

tiering

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

maxdata_on_san_application_components

The list of application components to be created.

Name	Type	Description
file_system	string	Defines the type of file system that will be installed on this application component.
host_management_url	string	The host management URL for this application component.
host_name	string	FQDN of the L2 host that contains the hot tier of this application component.
igroup_name	string	The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup.

Name	Type	Description
lun_count	integer	The number of LUNs in the application component.
metadata	array[metadata]	
name	string	The name of the application component.
protection_type	protection_type	
storage_service	storage_service	
tiering	maxdata_on_san_application_components_tiering	tiering
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]}

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application.
value	string	Value associated with the key.

maxdata_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

maxdata_on_san

MAX Data application using SAN.

Name	Type	Description
app_type	string	Type of the application that is being deployed on the L2.
application_components	array[maxdata_on_san_application_components]	The list of application components to be created.
metadata	array[metadata]	

Name	Type	Description
ocsm_url	string	The OnCommand System Manager URL for this application.
os_type	string	The name of the host OS running the application.

storage_service

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

dataset

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

mongo_db_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

protection_type

Name	Type	Description
local_rpo	string	The local RPO of the application.

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

secondary_igroups

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

mongo_db_on_san

MongoDB using SAN.

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

export_policy

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

component

Name	Type	Description
name	string	Name of the source component.

svm

Name	Type	Description
name	string	Name of the source SVM.

origin

Name	Type	Description
component	component	
svm	svm	

flexcache

Name	Type	Description
origin	origin	

nas_application_components_tiering

application-components.tiering

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

application_components

Name	Type	Description
export_policy	export_policy	
flexcache	flexcache	
name	string	The name of the application component.
qos	qos	
scale_out	boolean	Denotes a Flexgroup.
share_count	integer	The number of shares in the application component.
storage_service	storage_service	
tiering	nas_application_components_tiering	application-components.tiering
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]}

app_cifs_access

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

access.

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

app_nfs_access

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

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

protection_type

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

nas

A generic NAS application.

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

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

performance

Name	Type	Description
storage_service	storage_service	

hosts

Name	Type	Description
nqn	string	The host NQN.

zapp_nvme_components_subsystem

components.subsystem

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

zapp_nvme_components_tiering

application-components.tiering

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

components

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

local

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

rpo

Name	Type	Description
local	local	

zapp_nvme

An NVME application.

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

storage_service

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

archive_log

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

db

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

storage_service

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

ora_home

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

storage_service

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

redo_log

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

oracle_on_nfs

Oracle using NFS.

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

oracle_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

oracle_on_san

Oracle using SAN.

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

storage_service

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

grid_binary

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

storage_service

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

oracle_crs

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

oracle_rac_on_nfs

Oracle RAC using NFS.

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

db_sids

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

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

oracle_rac_on_san

Oracle RAC using SAN.

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

zapp_s3_bucket_application_components_access_policies_conditions

conditions

Name	Type	Description
delimiters	array[string]	
max_keys	array[integer]	
operator	string	Policy Condition Operator.

Name	Type	Description
prefixes	array[string]	
source_ips	array[string]	
usernames	array[string]	

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

Name	Type	Description
actions	array[string]	
conditions	array[zapp_s3_bucket_application_components_access_policies_conditions]	conditions.
effect	string	Allow or Deny Access.
principal	array[string]	
resources	array[string]	
sid	string	Statement Identifier.

exclude_aggregates

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

zapp_s3_bucket_application_components

The list of application components to be created.

Name	Type	Description
access_policies	array[zapp_s3_bucket_application_components_access_policies]	The list of S3 objectstore policies to be created.
capacity_tier	boolean	Prefer lower latency storage under similar media costs.

Name	Type	Description
comment	string	Object Store Server Bucket Description Usage: <(size 1..256)>
exclude_aggregates	array[exclude_aggregates]	
name	string	The name of the application component.
qos	qos	
size	integer	The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB MB GB TB PB]}
storage_service	storage_service	
uuid	string	Object Store Server Bucket UUID Usage: <UUID>

zapp_s3_bucket

A generic S3 bucket application.

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

san_application_components_tiering

application-components.tiering

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

application_components

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

san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

san

A generic SAN application.

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

storage_service

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

db

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

storage_service

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

log

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

sql_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

storage_service

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

temp_db

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

sql_on_san

Microsoft SQL using SAN.

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

access

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

sql_on_smb

Microsoft SQL using SMB.

Name	Type	Description
access	access	

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

storage_service

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

desktops

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

hyper_v_access

Name	Type	Description
service_account	string	Hyper-V service account.

vdi_on_nas

A VDI application using NAS.

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

vdi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

vdi_on_san

A VDI application using SAN.

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

storage_service

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

datastore

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

vsi_on_nas

A VSI application using NAS.

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

vsi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

vsi_on_san

A VSI application using SAN.

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

application_template

Application templates

Name	Type	Description
_links	self_link	

Name	Type	Description
cg_on_san	zapp_cg_on_san	A consistency group (CG) application using SAN.
description	string	Description.
maxdata_on_san	maxdata_on_san	MAX Data application using SAN.
missing_prerequisites	string	Missing prerequisites.
mongo_db_on_san	mongo_db_on_san	MongoDB using SAN.
name	string	Template name.
nas	nas	A generic NAS application.
nvme	zapp_nvme	An NVME application.
oracle_on_nfs	oracle_on_nfs	Oracle using NFS.
oracle_on_san	oracle_on_san	Oracle using SAN.
oracle_rac_on_nfs	oracle_rac_on_nfs	Oracle RAC using NFS.
oracle_rac_on_san	oracle_rac_on_san	Oracle RAC using SAN.
protocol	string	Access protocol.
s3_bucket	zapp_s3_bucket	A generic S3 bucket application.
san	san	A generic SAN application.
sql_on_san	sql_on_san	Microsoft SQL using SAN.
sql_on_smb	sql_on_smb	Microsoft SQL using SMB.
vdi_on_nas	vdi_on_nas	A VDI application using NAS.
vdi_on_san	vdi_on_san	A VDI application using SAN.
vsi_on_nas	vsi_on_nas	A VSI application using NAS.
vsi_on_san	vsi_on_san	A VSI application using SAN.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Retrieve an application template

GET /application/templates/{name}

Introduced In: 9.6

Retrieves an application template.

Template properties

Each application template has a set of properties. These properties are always nested under a property with the same name as the template. For example, when using the `mongo_db_on_san` template, the properties are found nested inside the `mongo_db_on_san` property. The properties nested under the template property are all specific to the template. The model for the application template object includes all the available templates, but only the object that corresponds to the template's name is returned, and only one is provided in any application API.

The model of each template includes a description of each property and its allowed values or usage. Default values are also indicated when available. The template properties returned by this API include an example value for each property.

Template prerequisites

Each template has a set of prerequisites required for its use. If any of these prerequisites are not met, the `missing_prerequisites` property indicates which prerequisite is missing.

Learn more

- [DOC /application](#)

Parameters

Name	Type	In	Required	Description
name	string	path	True	Template Name
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	self_link	
cg_on_san	zapp_cg_on_san	A consistency group (CG) application using SAN.
description	string	Description.
maxdata_on_san	maxdata_on_san	MAX Data application using SAN.
missing_prerequisites	string	Missing prerequisites.
mongo_db_on_san	mongo_db_on_san	MongoDB using SAN.
name	string	Template name.
nas	nas	A generic NAS application.
nvme	zapp_nvme	An NVME application.
oracle_on_nfs	oracle_on_nfs	Oracle using NFS.
oracle_on_san	oracle_on_san	Oracle using SAN.
oracle_rac_on_nfs	oracle_rac_on_nfs	Oracle RAC using NFS.
oracle_rac_on_san	oracle_rac_on_san	Oracle RAC using SAN.
protocol	string	Access protocol.
s3_bucket	zapp_s3_bucket	A generic S3 bucket application.

Name	Type	Description
san	san	A generic SAN application.
sql_on_san	sql_on_san	Microsoft SQL using SAN.
sql_on_smb	sql_on_smb	Microsoft SQL using SMB.
vdi_on_nas	vdi_on_nas	A VDI application using NAS.
vdi_on_san	vdi_on_san	A VDI application using SAN.
vsi_on_nas	vsi_on_nas	A VSI application using NAS.
vsi_on_san	vsi_on_san	A VSI application using SAN.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cg_on_san": {
    "cgs": [
      {
        "lun_containers": [
          {
            "luns": [
              {
                "igroups": [
                  {
                    "name": "string"
                  }
                ],
                "name": "string",
                "os_type": "string"
              }
            ],
            "name": "string"
          }
        ],
        "name": "string",
        "protection_type": {
          "local_policy": "string",
          "remote_rpo": "string"
        },
        "qos": {
          "policy": {
            "name": "string",
            "uuid": "string"
          }
        },
        "storage_service": {
          "name": "string"
        },
        "tiering": {
          "policy": "string"
        }
      }
    ]
  }
}
```

```

},
"description": "string",
"maxdata_on_san": {
  "app_type": "string",
  "application_components": [
    {}
  ],
  "metadata": [
    {
      "key": "string",
      "value": "string"
    }
  ],
  "ocsm_url": "string",
  "os_type": "string"
},
"missing_prerequisites": "string",
"name": "string",
"nas": {
  "application_components": [
    {}
  ],
  "cifs_access": [
    {
      "access": "string",
      "user_or_group": "string"
    }
  ],
  "nfs_access": [
    {
      "access": "string",
      "host": "string"
    }
  ],
  "protection_type": {
    "local_policy": "string",
    "local_rpo": "string",
    "remote_rpo": "string"
  }
},
"nvme": {
  "components": [
    {
      "name": "string",
      "os_type": "string",
      "performance": {

```

```

        "storage_service": {
            "name": "string"
        }
    },
    "qos": {
        "policy": {
            "name": "string",
            "uuid": "string"
        }
    },
    "subsystem": {
        "hosts": [
            {
                "nqn": "string"
            }
        ],
        "name": "string",
        "os_type": "string",
        "uuid": "string"
    },
    "tiering": {
        "policy": "string"
    }
}
],
"os_type": "string",
"rpo": {
    "local": {
        "name": "string",
        "policy": "string"
    }
}
},
"protocol": "string",
"s3_bucket": {
    "application_components": [
        {}
    ]
},
"san": {
    "application_components": [
        {
            "igroup_name": "string",
            "name": "string",
            "os_type": "string",
            "qos": {

```



```

        "policy": {
            "name": "string",
            "uuid": "string"
        },
        "storage_service": {
            "name": "string"
        },
        "tiering": {
            "policy": "string"
        }
    ],
    "os_type": "string",
    "protection_type": {
        "local_policy": "string",
        "local_rpo": "string",
        "remote_rpo": "string"
    }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

igroups

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

zapp_cg_on_san_cgs_lun_containers_luns

The list of LUNs to be created.

Name	Type	Description
count	integer	The number of LUNs.
igroups	array[igroups]	
name	string	The name/prefix of LUNs.
os_type	string	The name of the host OS.

zapp_cg_on_san_cgs_lun_containers

LUN containers.

Name	Type	Description
luns	array[zapp_cg_on_san_cgs_lun_containers_luns]	The list of LUNs to be created.

Name	Type	Description
name	string	The name or prefix of container (Volume). If the container exists, it becomes a part of the CG. If not, a new container is created.

protection_type

Name	Type	Description
local_policy	string	The Snapshot copy policy to apply to a CG.
remote_rpo	string	The remote RPO of the CG.

policy

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

qos

Name	Type	Description
policy	policy	

storage_service

Name	Type	Description
name	string	The storage service of the CG.

object_stores

zapp_cg_on_san_cgs_tiering

tiering

Name	Type	Description
policy	string	The storage tiering type of the CG.

zapp_cg_on_san_cgs

The list of CGs to be created.

Name	Type	Description
lun_containers	array[zapp_cg_on_san_cgs_lun_containers]	LUN containers.
name	string	The name of the consistency group.
protection_type	protection_type	
qos	qos	
storage_service	storage_service	
tiering	zapp_cg_on_san_cgs_tiering	tiering

zapp_cg_on_san

A consistency group (CG) application using SAN.

Name	Type	Description
cgs	array[zapp_cg_on_san_cgs]	The list of CGs to be created.

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application component.
value	string	Value associated with the key.

protection_type

Name	Type	Description
local_rpo	string	The local rpo of the application component.
remote_rpo	string	The remote rpo of the application component.

storage_service

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

maxdata_on_san_application_components_tiering

tiering

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

maxdata_on_san_application_components

The list of application components to be created.

Name	Type	Description
file_system	string	Defines the type of file system that will be installed on this application component.
host_management_url	string	The host management URL for this application component.
host_name	string	FQDN of the L2 host that contains the hot tier of this application component.
igroup_name	string	The name of the initiator group through which the contents of this application will be accessed. Modification of this parameter is a disruptive operation. All LUNs in the application component will be unmapped from the current igroup and re-mapped to the new igroup.
lun_count	integer	The number of LUNs in the application component.
metadata	array[metadata]	
name	string	The name of the application component.
protection_type	protection_type	

Name	Type	Description
storage_service	storage_service	
tiering	maxdata_on_san_application_components_tiering	tiering
total_size	integer	The total size of the application component, split across the member LUNs. Usage: {<integer>[KB MB GB TB PB]}

metadata

Name	Type	Description
key	string	Key to look up metadata associated with an application.
value	string	Value associated with the key.

maxdata_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

maxdata_on_san

MAX Data application using SAN.

Name	Type	Description
app_type	string	Type of the application that is being deployed on the L2.
application_components	array[maxdata_on_san_application_components]	The list of application components to be created.
metadata	array[metadata]	
ocsm_url	string	The OnCommand System Manager URL for this application.
os_type	string	The name of the host OS running the application.

storage_service

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

dataset

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

mongo_db_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

protection_type

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

secondary_igroups

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

mongo_db_on_san

MongoDB using SAN.

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

export_policy

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

component

Name	Type	Description
name	string	Name of the source component.

svm

Name	Type	Description
name	string	Name of the source SVM.

origin

Name	Type	Description
component	component	
svm	svm	

flexcache

Name	Type	Description
origin	origin	

nas_application_components_tiering

application-components.tiering

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

application_components

Name	Type	Description
export_policy	export_policy	
flexcache	flexcache	
name	string	The name of the application component.
qos	qos	
scale_out	boolean	Denotes a Flexgroup.
share_count	integer	The number of shares in the application component.
storage_service	storage_service	
tiering	nas_application_components_tiering	application-components.tiering
total_size	integer	The total size of the application component, split across the member shares. Usage: {<integer>[KB MB GB TB PB]}

app_cifs_access

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

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

app_nfs_access

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

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

protection_type

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

nas

A generic NAS application.

Name	Type	Description
application_components	array[application_components]	
cifs_access	array[app_cifs_access]	The list of CIFS access controls. You must provide either 'user_or_group' or 'access' to enable CIFS access.
nfs_access	array[app_nfs_access]	The list of NFS access controls. You must provide either 'host' or 'access' to enable NFS access.
protection_type	protection_type	

performance

Name	Type	Description
storage_service	storage_service	

hosts

Name	Type	Description
nqn	string	The host NQN.

zapp_nvme_components_subsystem

components.subsystem

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

zapp_nvme_components_tiering

application-components.tiering

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

components

Name	Type	Description
name	string	The name of the application component.
namespace_count	integer	The number of namespaces in the component.

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

local

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

rpo

Name	Type	Description
local	local	

zapp_nvme

An NVME application.

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

storage_service

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

archive_log

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

db

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

storage_service

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

ora_home

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

storage_service

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

redo_log

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

oracle_on_nfs

Oracle using NFS.

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

oracle_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

oracle_on_san

Oracle using SAN.

Name	Type	Description
archive_log	archive_log	
db	db	

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

storage_service

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

grid_binary

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

storage_service

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

oracle_crs

Name	Type	Description
copies	integer	The number of CRS volumes.

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

oracle_rac_on_nfs

Oracle RAC using NFS.

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

db_sids

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

oracle_rac_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

oracle_rac_on_san

Oracle RAC using SAN.

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

zapp_s3_bucket_application_components_access_policies_conditions

conditions

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

zapp_s3_bucket_application_components_access_policies

The list of S3 objectstore policies to be created.

Name	Type	Description
actions	array[string]	
conditions	array[zapp_s3_bucket_application_components_access_policies_conditions]	conditions.
effect	string	Allow or Deny Access.
principal	array[string]	

Name	Type	Description
resources	array[string]	
sid	string	Statement Identifier.

exclude_aggregates

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

zapp_s3_bucket_application_components

The list of application components to be created.

Name	Type	Description
access_policies	array[zapp_s3_bucket_application_components_access_policies]	The list of S3 objectstore policies to be created.
capacity_tier	boolean	Prefer lower latency storage under similar media costs.
comment	string	Object Store Server Bucket Description Usage: <(size 1..256)>
exclude_aggregates	array[exclude_aggregates]	
name	string	The name of the application component.
qos	qos	
size	integer	The total size of the S3 Bucket, split across the member components. Usage: {<integer>[KB MB GB TB PB]}
storage_service	storage_service	
uuid	string	Object Store Server Bucket UUID Usage: <UUID>

zapp_s3_bucket

A generic S3 bucket application.

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

san_application_components_tiering

application-components.tiering

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

application_components

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

san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

san

A generic SAN application.

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

storage_service

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

db

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

storage_service

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

log

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

sql_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

storage_service

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

temp_db

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

sql_on_san

Microsoft SQL using SAN.

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

access

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

sql_on_smb

Microsoft SQL using SMB.

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

storage_service

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

desktops

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

hyper_v_access

Name	Type	Description
service_account	string	Hyper-V service account.

vdi_on_nas

A VDI application using NAS.

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

vdi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

vdi_on_san

A VDI application using SAN.

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

storage_service

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

datastore

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

vsi_on_nas

A VSI application using NAS.

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

vsi_on_san_new_igroups

The list of initiator groups to create.

Name	Type	Description
initiators	array[string]	

vsi_on_san

A VSI application using SAN.

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

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Cloud

Cloud overview

Manages cloud (object storage) targets

Manage cloud targets

Cloud targets endpoint overview

Retrieving a collection of cloud targets

The cloud targets GET API retrieves all cloud targets defined in the cluster.

Creating cloud targets

The cluster administrator tells ONTAP how to connect to a cloud target. The following pre-requisites must be met before creating an object store configuration in ONTAP.

A valid data bucket or container must be created with the object store provider. This assumes that the user has valid account credentials with the object store provider to access the data bucket.

The ONTAP node must be able to connect to the object store.

This includes:

- Fast, reliable connectivity to the object store.
- An inter-cluster LIF (logical interface) must be configured on the cluster. ONTAP verifies connectivity prior to saving this configuration information.
- If SSL/TLS authentication is required, then valid certificates must be installed.
- FabricPool license (required for all object stores except SGWS).

Deleting cloud targets

If a cloud target is used by an aggregate, then the aggregate must be deleted before the cloud target can be deleted.

Retrieve cloud targets in the cluster

GET /cloud/targets

Introduced In: 9.6

Retrieves the collection of cloud targets in the cluster.

Related ONTAP commands

- `storage aggregate object-store config show`

Parameters

Name	Type	In	Required	Description
port	integer	query	False	Filter by port
azure_account	string	query	False	Filter by azure_account
snapmirror_use	string	query	False	Filter by snapmirror_use
ssl_enabled	boolean	query	False	Filter by ssl_enabled
access_key	string	query	False	Filter by access_key

Name	Type	In	Required	Description
certificate_validation_enabled	boolean	query	False	Filter by certificate_validation_enabled
uuid	string	query	False	Filter by uuid
owner	string	query	False	Filter by owner
name	string	query	False	Filter by name
use_http_proxy	boolean	query	False	Filter by use_http_proxy <ul style="list-style-type: none"> Introduced in: 9.7
server_side_encryption	string	query	False	Filter by server_side_encryption <ul style="list-style-type: none"> Introduced in: 9.7
server	string	query	False	Filter by server
provider_type	string	query	False	Filter by provider_type
cap_url	string	query	False	Filter by cap_url
cluster.uuid	string	query	False	Filter by cluster.uuid <ul style="list-style-type: none"> Introduced in: 9.7
cluster.name	string	query	False	Filter by cluster.name <ul style="list-style-type: none"> Introduced in: 9.7
container	string	query	False	Filter by container
ipspace.name	string	query	False	Filter by ipspace.name

Name	Type	In	Required	Description
ipspace.uuid	string	query	False	Filter by ipspace.uuid
authentication_type	string	query	False	Filter by authentication_type
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
url_style	string	query	False	Filter by url_style • Introduced in: 9.8
used	integer	query	False	Filter by used
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cloud_target]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access_key": "string",
      "authentication_type": "string",
      "azure_account": "string",
      "cap_url":
"https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&mission=mymission&role=myrole",
      "cluster": {
        "name": "string",
        "uuid": "string"
      },
      "container": "bucket1",
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "exchange",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "string",
      "owner": "string",
      "provider_type": "string",
      "server": "string",
      "server_side_encryption": "string",
      "snapmirror_use": "string",
      "svm": {
        "_links": {
```

```

        "self": {
            "href": "/api/resourcelink"
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "url_style": "string",
    "used": 0,
    "uuid": "string"
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

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

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
name	string	The name of the cluster that owns the cloud target. For POST, this accepts the name of the peer cluster only if the cluster is in switchover state.
uuid	string	The UUID of the cluster that owns the cloud target. For POST, this accepts the UUID of the peer cluster only if the cluster is in switchover state.

ipspace

IPspace to use in order to reach the cloud target.

Name	Type	Description
_links	_links	
name	string	IPspace name
uuid	string	IPspace UUID

svm

This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is

required for SnapMirror targets and not allowed for FabricPool targets.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cloud_target

Name	Type	Description
_links	_links	
access_key	string	Access key ID for AWS_S3 and other S3 compatible provider types.
authentication_type	string	Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST.
azure_account	string	Azure account
cap_url	string	This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store.
certificate_validation_enabled	boolean	Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types. <ul style="list-style-type: none">• Introduced in: 9.6
cluster	cluster	
container	string	Data bucket/container name
ipspace	ipspace	IPspace to use in order to reach the cloud target.
name	string	Cloud target name

Name	Type	Description
owner	string	Owner of the target. Allowed values are FabricPool or SnapMirror. A target can be used by only one feature.
port	integer	Port number of the object store that ONTAP uses when establishing a connection. Required in POST. <ul style="list-style-type: none"> Introduced in: 9.6
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, SGWS, and ONTAP_S3 are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS. <ul style="list-style-type: none"> Introduced in: 9.6 readCreate: 1
server	string	Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3- <region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region>

Name	Type	Description
server_side_encryption	string	<p>Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. This is an advanced property. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.</p> <ul style="list-style-type: none"> enum: ["none", "sse_s3"] Introduced in: 9.7
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	svm	This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.
url_style	string	URL style used to access S3 bucket.
use_http_proxy	boolean	Use HTTP proxy when connecting to the object store.
used	integer	The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes.
uuid	string	Cloud target UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Create a cloud target

POST /cloud/targets

Introduced In: 9.6

Creates a cloud target.

Required properties

- `name` - Name for the cloud target.
- `owner` - Owner of the target: *fabricpool*, *snapmirror*.
- `provider_type` - Type of cloud provider: *AWS_S3*, *Azure_Cloud*, *SGWS*, *IBM_COS*, *AliCloud*, *GoogleCloud*, *ONTAP_S3*.
- `server` - Fully qualified domain name of the object store server. Required when `provider_type` is one of the following: *SGWS*, *IBM_COS*, *AliCloud*.
- `container` - Data bucket/container name.
- `access_key` - Access key ID if `provider_type` is not *Azure_Cloud* and `authentication_type` is *key*.
- `secret_password` - Secret access key if `provider_type` is not *Azure_Cloud* and `authentication_type` is *key*.
- `azure_account` - Azure account if `provider_type` is *Azure_Cloud*.
- `azure_private_key` - Azure access key if `provider_type` is *Azure_Cloud*.
- `cap_url` - Full URL of the request to a CAP server for retrieving temporary credentials if `authentication_type` is *cap*.

- `svm.name` or `svm.uuid` - Name or UUID of SVM if owner is *snapmirror*.
- `snapmirror_use` - Use of the cloud target if owner is *snapmirror*: data, metadata.

Recommended optional properties

- `authentication_type` - Authentication used to access the target: *key*, *cap*, *ec2_iam*, *gcp_sa*, *azure_msi*.
- `ssl_enabled` - SSL/HTTPS enabled or disabled.
- `port` - Port number of the object store that ONTAP uses when establishing a connection.
- `ipspace` - IPspace to use in order to reach the cloud target.
- `use_http_proxy` - Use the HTTP proxy when connecting to the object store server.

Default property values

- `authentication_type`
- *ec2_iam* - if running in Cloud Volumes ONTAP in AWS
- *gcp_sa* - if running in Cloud Volumes ONTAP in GCP
- *azure_msi* - if running in Cloud Volumes ONTAP in Azure
- *key* - in all other cases.
- `server`
- *s3.amazonaws.com* - if `provider_type` is *AWS_S3*
- *blob.core.windows.net* - if `provider_type` is *Azure_Cloud*
- *storage.googleapis.com* - if `provider_type` is *GoogleCloud*
- `ssl_enabled` - *true*
- `port`
- *443* if `ssl_enabled` is *true* and `provider_type` is not *SGWS*
- *8082* if `ssl_enabled` is *true* and `provider_type` is *SGWS*
- *80* if `ssl_enabled` is *false* and `provider_type` is not *SGWS*
- *8084* if `ssl_enabled` is *false* and `provider_type` is *SGWS*
- `ipspace` - *Default*
- `certificate_validation_enabled` - *true*
- `ignore_warnings` - *false*
- `check_only` - *false*
- `use_http_proxy` - *false*
- `server_side_encryption`
- *none* - if `provider_type` is *ONTAP_S3*
- *sse_s3* - if `provider_type` is not *ONTAP_S3*

- `url_style`
- `path_style` - if `provider_type` is neither `AWS_S3` nor `AliCloud`
- `virtual_hosted_style` - if `provider_type` is either `AWS_S3` or `AliCloud`

Related ONTAP commands

- `storage aggregate object-store config create`

Parameters

Name	Type	In	Required	Description
<code>ignore_warnings</code>	boolean	query	False	Specifies whether or not warning codes should be ignored.
<code>check_only</code>	boolean	query	False	Do not create the target configuration, only check that the POST request succeeds.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
access_key	string	Access key ID for AWS_S3 and other S3 compatible provider types.

Name	Type	Description
authentication_type	string	Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST.
azure_account	string	Azure account
azure_private_key	string	Azure access key
cap_url	string	This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store.
certificate_validation_enabled	boolean	Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types. • Introduced in: 9.6
cluster	cluster	
container	string	Data bucket/container name
ipspace	ipspace	IPspace to use in order to reach the cloud target.
name	string	Cloud target name
owner	string	Owner of the target. Allowed values are FabricPool or SnapMirror. A target can be used by only one feature.
port	integer	Port number of the object store that ONTAP uses when establishing a connection. Required in POST. • Introduced in: 9.6

Name	Type	Description
provider_type	string	<p>Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, SGWS, and ONTAP_S3 are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS.</p> <ul style="list-style-type: none"> • Introduced in: 9.6 • readCreate: 1
secret_password	string	Secret access key for AWS_S3 and other S3 compatible provider types.
server	string	<p>Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region></p>

Name	Type	Description
server_side_encryption	string	<p>Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. This is an advanced property. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.</p> <ul style="list-style-type: none"> enum: ["none", "sse_s3"] Introduced in: 9.7
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	svm	This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.
url_style	string	URL style used to access S3 bucket.
use_http_proxy	boolean	Use HTTP proxy when connecting to the object store.
used	integer	The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes.
uuid	string	Cloud target UUID

Example request

```
{
  "access_key": "string",
  "authentication_type": "string",
  "azure_account": "string",
  "azure_private_key": "string",
  "cap_url":
  "https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&mission=mymission&role=myrole",
  "cluster": {
    "name": "string",
    "uuid": "string"
  },
  "container": "bucket1",
  "ipspace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "string",
  "owner": "string",
  "provider_type": "string",
  "secret_password": "string",
  "server": "string",
  "server_side_encryption": "string",
  "snapmirror_use": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "url_style": "string",
  "used": 0,
  "uuid": "string"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster

Name	Type	Description
name	string	The name of the cluster that owns the cloud target. For POST, this accepts the name of the peer cluster only if the cluster is in switchover state.
uuid	string	The UUID of the cluster that owns the cloud target. For POST, this accepts the UUID of the peer cluster only if the cluster is in switchover state.

ipspace

IPspace to use in order to reach the cloud target.

Name	Type	Description
name	string	IPspace name
uuid	string	IPspace UUID

svm

This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cloud_target

Name	Type	Description
access_key	string	Access key ID for AWS_S3 and other S3 compatible provider types.
authentication_type	string	Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST.
azure_account	string	Azure account
azure_private_key	string	Azure access key
cap_url	string	This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-pasword, and session token) for accessing the object store.
certificate_validation_enabled	boolean	Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types. • Introduced in: 9.6
cluster	cluster	
container	string	Data bucket/container name
ipspace	ipspace	IPspace to use in order to reach the cloud target.
name	string	Cloud target name
owner	string	Owner of the target. Allowed values are FabricPool or SnapMirror. A target can be used by only one feature.

Name	Type	Description
port	integer	Port number of the object store that ONTAP uses when establishing a connection. Required in POST. <ul style="list-style-type: none"> Introduced in: 9.6
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, SGWS, and ONTAP_S3 are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS. <ul style="list-style-type: none"> Introduced in: 9.6 readCreate: 1
secret_password	string	Secret access key for AWS_S3 and other S3 compatible provider types.
server	string	Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region>

Name	Type	Description
server_side_encryption	string	<p>Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. This is an advanced property. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.</p> <ul style="list-style-type: none"> enum: ["none", "sse_s3"] Introduced in: 9.7
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	svm	This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.
url_style	string	URL style used to access S3 bucket.
use_http_proxy	boolean	Use HTTP proxy when connecting to the object store.
used	integer	The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes.
uuid	string	Cloud target UUID

job_link

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

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

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

Delete a cloud target

DELETE /cloud/targets/{uuid}

Introduced In: 9.6

Deletes the cloud target specified by the UUID. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate object-store config delete`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cloud target UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

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

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

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

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a cloud target

GET /cloud/targets/{uuid}

Introduced In: 9.6

Retrieves the cloud target specified by the UUID.

Related ONTAP commands

- `storage aggregate object-store config show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cloud target UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
access_key	string	Access key ID for AWS_S3 and other S3 compatible provider types.
authentication_type	string	Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST.
azure_account	string	Azure account
cap_url	string	This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store.
certificate_validation_enabled	boolean	Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types. <ul style="list-style-type: none">• Introduced in: 9.6

Name	Type	Description
cluster	cluster	
container	string	Data bucket/container name
ipspace	ipspace	IPspace to use in order to reach the cloud target.
name	string	Cloud target name
owner	string	Owner of the target. Allowed values are FabricPool or SnapMirror. A target can be used by only one feature.
port	integer	Port number of the object store that ONTAP uses when establishing a connection. Required in POST. <ul style="list-style-type: none"> Introduced in: 9.6
provider_type	string	Type of cloud provider. Allowed values depend on owner type. For FabricPool, AliCloud, AWS_S3, Azure_Cloud, GoggleCloud, IBM_COS, SGWS, and ONTAP_S3 are allowed. For SnapMirror, the valid values are AWS_S3 or SGWS. <ul style="list-style-type: none"> Introduced in: 9.6 readCreate: 1
server	string	Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region>

Name	Type	Description
server_side_encryption	string	<p>Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. This is an advanced property. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.</p> <ul style="list-style-type: none"> enum: ["none", "sse_s3"] Introduced in: 9.7
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	svm	This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.
url_style	string	URL style used to access S3 bucket.
use_http_proxy	boolean	Use HTTP proxy when connecting to the object store.
used	integer	The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes.
uuid	string	Cloud target UUID

Example response

```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access_key": "string",
  "authentication_type": "string",
  "azure_account": "string",
  "cap_url":
  "https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&mission=mymission&role=myrole",
  "cluster": {
    "name": "string",
    "uuid": "string"
  },
  "container": "bucket1",
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "string",
  "owner": "string",
  "provider_type": "string",
  "server": "string",
  "server_side_encryption": "string",
  "snapmirror_use": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "url_style": "string",
  "used": 0,
  "uuid": "string"
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
name	string	The name of the cluster that owns the cloud target. For POST, this accepts the name of the peer cluster only if the cluster is in switchover state.
uuid	string	The UUID of the cluster that owns the cloud target. For POST, this accepts the UUID of the peer cluster only if the cluster is in switchover state.

ipspace

IPspace to use in order to reach the cloud target.

Name	Type	Description
_links	_links	
name	string	IPspace name
uuid	string	IPspace UUID

svm

This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a cloud target

PATCH /cloud/targets/{uuid}

Introduced In: 9.6

Updates the cloud target specified by the UUID with the fields in the body. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate object-store config modify`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cloud target UUID
ignore_warnings	boolean	query	False	Specifies whether or not warnings should be ignored.

Name	Type	In	Required	Description
check_only	boolean	query	False	Do not modify the configuration, only check that the PATCH request succeeds.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
access_key	string	Access key ID for AWS_S3 and other S3 compatible provider types.
authentication_type	string	Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST.

Name	Type	Description
azure_account	string	Azure account
azure_private_key	string	Azure access key
cap_url	string	This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store.
certificate_validation_enabled	boolean	Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types. <ul style="list-style-type: none"> Introduced in: 9.6
cluster	cluster	
name	string	Cloud target name
port	integer	Port number of the object store that ONTAP uses when establishing a connection. Required in POST. <ul style="list-style-type: none"> Introduced in: 9.6
secret_password	string	Secret access key for AWS_S3 and other S3 compatible provider types.
server	string	Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region>

Name	Type	Description
server_side_encryption	string	<p>Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. This is an advanced property. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.</p> <ul style="list-style-type: none"> enum: ["none", "sse_s3"] Introduced in: 9.7
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	svm	This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.
url_style	string	URL style used to access S3 bucket.
use_http_proxy	boolean	Use HTTP proxy when connecting to the object store.
used	integer	The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes.
uuid	string	Cloud target UUID

Example request

```
{
  "access_key": "string",
  "authentication_type": "string",
  "azure_account": "string",
  "azure_private_key": "string",
  "cap_url":
    "https://123.45.67.89:1234/CAP/api/v1/credentials?agency=myagency&mission=mymission&role=myrole",
  "name": "string",
  "secret_password": "string",
  "server": "string",
  "server_side_encryption": "string",
  "snapmirror_use": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "url_style": "string",
  "used": 0,
  "uuid": "string"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster

ipspace

IPspace to use in order to reach the cloud target.

Name	Type	Description
name	string	IPspace name
uuid	string	IPspace UUID

svm

This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cloud_target

Name	Type	Description
access_key	string	Access key ID for AWS_S3 and other S3 compatible provider types.
authentication_type	string	Authentication used to access the target. SnapMirror does not yet support CAP. Required in POST.
azure_account	string	Azure account
azure_private_key	string	Azure access key

Name	Type	Description
cap_url	string	This parameter is available only when auth-type is CAP. It specifies a full URL of the request to a CAP server for retrieving temporary credentials (access-key, secret-password, and session token) for accessing the object store.
certificate_validation_enabled	boolean	Is SSL/TLS certificate validation enabled? The default value is true. This can only be modified for SGWS, IBM_COS, and ONTAP_S3 provider types. • Introduced in: 9.6
cluster	cluster	
name	string	Cloud target name
port	integer	Port number of the object store that ONTAP uses when establishing a connection. Required in POST. • Introduced in: 9.6
secret_password	string	Secret access key for AWS_S3 and other S3 compatible provider types.
server	string	Fully qualified domain name of the object store server. Required on POST. For Amazon S3, server name must be an AWS regional endpoint in the format s3.amazonaws.com or s3-<region>.amazonaws.com, for example, s3-us-west-2.amazonaws.com. The region of the server and the bucket must match. For Azure, if the server is a "blob.core.windows.net" or a "blob.core.usgovcloudapi.net", then a value of azure-account followed by a period is added in front of the server.</region>

Name	Type	Description
server_side_encryption	string	<p>Encryption of data at rest by the object store server for AWS_S3 and other S3 compatible provider types. This is an advanced property. In most cases it is best not to change default value of "sse_s3" for object store servers which support SSE-S3 encryption. The encryption is in addition to any encryption done by ONTAP at a volume or at an aggregate level. Note that changing this option does not change encryption of data which already exist in the object store.</p> <ul style="list-style-type: none"> enum: ["none", "sse_s3"] Introduced in: 9.7
snapmirror_use	string	Use of the cloud target by SnapMirror.
ssl_enabled	boolean	SSL/HTTPS enabled or not
svm	svm	This field is only applicable when used for SnapMirror. For POST and PATCH, SVM information is required for SnapMirror targets and not allowed for FabricPool targets.
url_style	string	URL style used to access S3 bucket.
use_http_proxy	boolean	Use HTTP proxy when connecting to the object store.
used	integer	The amount of cloud space used by all the aggregates attached to the target, in bytes. This field is only populated for FabricPool targets. The value is recalculated once every 5 minutes.
uuid	string	Cloud target UUID

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Cluster

Cluster overview

Overview

These APIs enable you to perform a number of independent workflows, including:

- Creating the cluster
- Adding nodes to the cluster
- Managing cluster configuration data (including name, version, NTP servers, name servers, and DNS domains)
- Managing node configuration data (including node names, models, serial numbers, and HA group information)
- Discovering the nodes on the cluster network that can be added to the cluster
- Viewing and updating current and recent jobs
- Updating the cluster software

Pre-Cluster APIs

A few of the cluster APIs (namely, POST/OPTIONS on /api/cluster, GET/HEAD/OPTIONS on /api/cluster/nodes, and calls on /api/cluster/jobs) are allowed before the cluster is created. These APIs support creation of the cluster and monitoring of its progress. Any other cluster API used before the cluster is created will fail.

Manage clusters

Cluster endpoint overview

Overview

You can use this API to create a cluster, update cluster-wide configurations, and retrieve the current configuration details.

Creating a cluster

You can create a new cluster by issuing a POST request to /cluster. Parameters are provided in the body of the POST request to configure cluster-wide settings and add nodes during the cluster setup.

Fields used for creating a cluster

The fields used for the cluster APIs fall into the following categories:

- Required cluster-wide configuration
- Optional cluster-wide configuration

Required cluster-wide configuration

The following fields are always required for any POST /cluster request:

- name
- password

Optional cluster-wide configuration

The following fields are used to set up additional cluster-wide configurations:

- location
- contact
- dns_domains
- name_servers
- ntp_servers
- timezone
- license
- configuration_backup
- management_interface
- nodes

Nodes field

The nodes field specifies the nodes to join to the cluster. To use this API, all nodes must run the same version of ONTAP. If you do not specify a node, the cluster is configured with one node added. The REST request is issued to the node that is added to the cluster. If you specify one node, do not use the "node.cluster_interface.ip.address" field. If you specify multiple nodes, specify the node to which the REST request is issued in addition to the remote nodes. Use the "node.cluster_interface.ip.address" field to identify each node. All other node fields are optional in all cases. If you provide a field for one node, you need to provide the same field for all nodes.

Node networking fields

The cluster management interface and each node management interface use the cluster management interface subnet mask and gateway. For advanced configurations in which the cluster and node management interfaces are on different subnets, use the /network/ip/interface APIs to configure network interfaces after setup is complete.

The management interfaces are used to communicate with the name servers and NTP servers. The address family of the name servers and NTP servers must match the management interfaces address family.

Single node cluster field

When the "single_node_cluster" field is set to "true", the cluster is created in single node cluster mode. You can provide a node field for this node for node-specific configuration but do not use the "node.cluster_interface.ip.address" field. Storage failover is configured to non-HA mode, and ports used for cluster ports are moved to the default IPspace. This might cause the node to reboot during setup. While a node reboots, the RESTful interface might not be available. See "Connection failures during cluster create" for more information.

Create recommended aggregates parameter

When the "create_recommended_aggregates" parameter is set to "true", aggregates based on an optimal layout recommended by the system are created on each of the nodes in the cluster. The default setting is "false".

Performance monitoring

Performance of the cluster can be monitored by the `metric.*` and `statistics.*` fields. These fields show the performance of the cluster in terms of IOPS, latency and throughput. The `metric.*` fields denote an average, whereas the `statistics.*` fields denote a real-time monotonically increasing value aggregated across all nodes.

Monitoring cluster create status

Errors before the job starts

Configuration in the POST /cluster request is validated before the cluster create job starts. If an invalid configuration is found, an HTTP error code in the 4xx range is returned. No cluster create job is started.

Polling on the job

After a successful POST /cluster request is issued, an HTTP error code of 202 is returned along with a job UUID and link in the body of the response. The cluster create job continues asynchronously and is monitored with the job UUID using the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

Errors during the job

If a failure occurs during the cluster create job, the job body provides details of the error along with error code fields. See the error table under "Responses" in the POST /cluster documentation for common error codes and descriptions.

Rerunning POST /cluster

The POST /cluster request can be rerun if errors occur. When rerunning the request, use the same body and query parameters. You can change the value of any field in the original body or query, but you cannot change the provided fields. For example, an initial request might have a body section as follows:

+

```
body =
{
  "name": "clusCreateRerun",
  "password": "openSesame",
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "1.1.1.1"
        }
      }
    },
    {
      "cluster_interface": {
        "ip": {
          "address": "2.2.2.2"
        }
      }
    }
  ]
}
```

A rerun request updates the body details to:

+

```
body =
{
  "name": "clusCreateRerun",
  "password": "openSesame",
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "3.3.3.3"
        }
      }
    },
    {
      "cluster_interface": {
        "ip": {
          "address": "4.4.4.4"
        }
      }
    }
  ]
}
```

A rerun request with the following body details is invalid:

+

```
body =
{
  "name": "clusCreateRerun",
  "password": "openSesame",
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "3.3.3.3"
        }
      }
    }
  ]
}
```

Note that the password might already be configured. If a password is already configured and then a new password is provided, the new request overwrites the existing password. If a password is already configured either by another interface or by a previous POST request to /cluster, authenticate any future REST requests with that password. If a POST request to /cluster with the default return_timeout of 0 returns an error, then the password was not changed.

Connection failures during cluster create

A request to poll the job status might fail during a cluster create job in the following two cases. In these cases, programmatic use of the RESTful interface might be resilient to these connection failures.

1. When the "single_node_cluster" flag is set to "true", the node might reboot. During this time, the RESTful interface might refuse connections and return errors on a GET request, or connection timeouts might occur. Programmatic use of the RESTful interface during reboots must consider these effects while polling a cluster create job.
2. The "mgmt_auto" LIF is removed during the cluster create job. A POST /cluster request might be issued on the "mgmt_auto" LIF. However, requests to poll the job status might fail during cluster create when the "mgmt_auto" LIF is removed. The "mgmt_auto" LIF is only removed if a cluster management interface is provided as an argument to POST /cluster, and only after the cluster management interface is created. Programmatic use of the POST /cluster API on the "mgmt_auto" LIF should be configured to dynamically switch to polling the job on the cluster management LIF.

Modifying cluster configurations

The following fields can be used to modify a cluster-wide configuration:

- name
- location
- contact
- dns_domains
- name_servers
- timezone
- certificate

Examples

Minimally configuring a 2-node setup

```
# Body
minimal_2_node_cluster.txt(body):
{
  "name": "clusCreateExample1",
  "password": "openSesame",
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "1.1.1.1"
        }
      }
    },
    {
      "cluster_interface": {
        "ip": {
          "address": "2.2.2.2"
        }
      }
    }
  ]
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster" -d
"@minimal_2_node_cluster.txt"
```

Setting up a single node with additional node configuration and auto aggregate creation

```
# Body
single_node_additional_config.txt (body):
{
  "name": "clusCreateExample2",
  "password": "openSesame",
  "nodes": [
    {
      "name": "singleNode",
      "location": "Sunnyvale"
    }
  ]
}

# Request
curl -X POST "https://<mgmt-
ip>/api/cluster?single_node_cluster=true&create_recommended_aggregates=true" -d "@single_node_additional_config.txt"
```

Modifying a cluster-wide configuration

```
# Body
modify_cluster_config.txt (body):
{
  "contact": "it@company.com"
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster" -d
"@modify_cluster_config.txt"
```

Creating a cluster using the cluster "create" operation

This example shows how to create a cluster using the cluster APIs. Specifically, this example shows the creation of a two-node cluster and uses information from the nodes themselves combined with user supplied information to configure the cluster.

Preparing for setup

Before the REST APIs can be issued to create the cluster, the cluster must be wired up and powered on. The network connections between the nodes for the cluster interconnect and the connections to the management network must be completed. After the nodes are powered on, the nodes automatically configure interfaces on the platform's default cluster ports to allow the nodes to discover each other during setup and expansion workflows. You must configure a management interface on one node or use the `mgmt_auto` LIF, which is

assigned an IP address using DHCP, to start using the REST APIs. By making a console connection to a node, the cluster setup wizard guides you through the configuration of the initial node management interface to which the REST calls can be sent. Once this step is completed, exit the wizard by typing "exit". You can then issue REST API requests.

1. Wire and power on the nodes.
2. Make a console connection to one node to access the cluster setup wizard.
3. Enter node management interface information to enable REST API requests to be sent to the node.

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
  http://support.netapp.com/autosupport/
Type yes to confirm and continue {yes}: yes
Enter the node management interface port [e0c]:
  Enter the node management interface IP address: 10.224.82.249
  Enter the node management interface netmask: 255.255.192.0
  Enter the node management interface default gateway: 10.224.64.1
  A node management interface on port e0c with IP address 10.224.82.249
has been created.
  Use your web browser to complete cluster setup by accessing
  https://10.224.82.249
  Otherwise, press Enter to complete cluster setup using the command
line
  interface: exit
  Exiting the cluster setup wizard. Any changes you made have been
saved.
  The cluster administrator's account (username "admin") password is set
to the system default.
  Warning: You have exited the cluster setup wizard before completing
all
  of the tasks. The cluster is not configured. You can complete cluster
setup by typing
  "cluster setup" in the command line interface.
```

Discovering the nodes

If you issue a `GET /api/cluster/nodes` request when the nodes are not in a cluster, the API returns a list of nodes that were discovered on the cluster interconnect. Information returned includes the node's serial number, model, software version, UUID, and cluster interface address. The number of nodes returned should

be the same as the number of nodes expected to be in the cluster. If too many nodes are discovered, remove the nodes that should not be part of the cluster. If not enough nodes are discovered, verify all the nodes are powered on, that the connections to the cluster interconnect are complete, and retry the command.

+

```
# The API:
/api/cluster/nodes

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/nodes?fields=state,uptime" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "6dce4710-c860-11e9-b5bc-005056bb6135",
      "name": "cluster1",
      "uptime": 134555,
      "state": "up",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/6dce4710-c860-11e9-b5bc-005056bb6135"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/cluster/nodes?fields=state,uptime"
    }
  }
}
```

Creating the cluster

When the node information is available, including each node's cluster interface address, you can assemble the information for creating the cluster. Provide the cluster name and the password for the admin account. The rest of the information is optional and can be configured later using other APIs. Provide the cluster interface address for each node to be included in the cluster so that you can connect to it while adding it to the cluster. In addition to the cluster interface address, you can provide the optional node name, location, and management interface information. If you do not provide node names, nodes are named based on the cluster name. The nodes' management interface subnet mask and gateway values are omitted and must be the same as the cluster management interface's subnet mask and gateway.

+


```
# The API:
/api/cluster

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster" -H "accept:
application/hal+json" -H "accept: application/hal+json" -d
'{"name":"cluster1","location":"datacenter1","contact":"me","dns_domains":
["example.com"],"name_servers":["10.224.223.130","10.224.223.131","10.224.
223.132"],"ntp_servers":["time.nist.gov"],"management_interface":{"ip":{"a
ddress":"10.224.82.25","netmask":"255.255.192.0","gateway":"10.224.64.1"}}
,"password":"mypassword","license":{"keys":["AMEPOSIOIKLKGEEEEEDGNDEKSJDE"]
},"nodes":[{"cluster_interface":{"ip":{"address":"169.254.245.113"}}, {"name"
:"node1","management_interface":{"ip":{"address":"10.224.82.29"}}}, {"clust
er_interface":{"ip":{"address":"169.254.217.95"}}, {"name":"node2","manageme
nt_interface":{"ip":{"address":"10.224.82.31"}}}]}'

# The response:
{
  "job": {
    "uuid": "b5bc07e2-19e9-11e9-a751-005056bbd95f",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f"
      }
    }
  }
}
}
```

Monitoring the progress of cluster creation

To monitor the progress of the cluster create operation, poll the returned job link until the state value is no longer "running" or "queued".

+

```
# The API:
/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/b5bc07e2-1e9-11e9-a751-005056bbd95f" -H "accept: application/hal+json"

# The response:
{
  "uuid": "b5bc07e2-19e9-11e9-a751-005056bbd95f",
  "description": "POST /api/cluster",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b5bc07e2-19e9-11e9-a751-005056bbd95f"
    }
  }
}
```

Verifying the cluster information

After the cluster is created, you can verify the information applied using a number of APIs. You can retrieve most of the information provided using the `/api/cluster` and `/api/cluster/nodes` APIs. In addition, you can view the network interface and route information using the `/api/network` APIs. The following example shows how to retrieve the cluster information:

+

```
# The API:
/api/cluster

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster?fields=management_interfaces"
-H "accept: application/hal+json"

# The response:
{
  "management_interfaces": [
    {
      "uuid": "c661725a-19e9-11e9-a751-005056bbd95f",
      "name": "cluster_mgmt",
      "ip": {
        "address": "10.224.82.25"
      },
      "_links": {
        "self": {
          "href": "/api/network/ip/interfaces/c661725a-19e9-11e9-a751-005056bbd95f"
        }
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster"
    }
  }
}
```

Retrieve a cluster configuration

GET /cluster

Introduced In: 9.6

Retrieves the cluster configuration.

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
certificate	certificate	Certificate used by cluster and node management interfaces for TLS connection requests.
contact	string	

Name	Type	Description
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
location	string	
management_interfaces	array[management_interfaces]	
metric	metric	Performance numbers, such as IOPS latency and throughput.
name	string	
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
ntp_servers	array[string]	Host name, IPv4 address, or IPv6 address for the external NTP time servers.
san_optimized	boolean	Specifies if this cluster is an All SAN Array.

Name	Type	Description
statistics	statistics	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
timezone	timezone	<p>Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:</p> <ul style="list-style-type: none"> • console messages; • logging to internal ONTAP log files; and • localized REST API full ISO-8601 date, time, and time zone format information. <p>Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.</p> <ul style="list-style-type: none"> • Introduced in: 9.7
uuid	string	
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "location": "building 1",
  "management_interfaces": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
```

```

    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"name": "cluster1",
"name_servers": [
  "10.224.65.20",
  "2001:db08:a0b:12f0::1"
],
"ntp_servers": [
  "time.nist.gov",
  "10.98.19.20",
  "2610:20:6F15:15::27"
],
"statistics": {
  "iops_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},

```



```

"timezone": {
  "name": "America/New_York"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"version": {
  "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
  "generation": "9",
  "major": "4",
  "minor": "0"
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

certificate

Certificate used by cluster and node management interfaces for TLS connection requests.

Name	Type	Description
_links	_links	
name	string	Certificate name
uuid	string	Certificate UUID

configuration_backup

Name	Type	Description
password	string	
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	
validate_certificate	boolean	Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default.

license

License keys or NLF contents.

Name	Type	Description
keys	array[string]	

ip

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

management_interface

The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	ip	Object to setup an interface along with its default router.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

management_interfaces

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.

Name	Type	Description
uuid	string	The UUID that uniquely identifies the interface.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	List of FRUs on the node. Only returned when requested by name.

Name	Type	Description
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

failure

Indicates the failure code and message.

Name	Type	Description
code	integer	Message code
message	string	Detailed message based on the state.

giveback

Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ports

Name	Type	Description
number	integer	HA port number

Name	Type	Description
state	string	<p>HA port state:</p> <ul style="list-style-type: none"> • <i>down</i> - Logical HA link is down. • <i>initialized</i> - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port. • <i>armed</i> - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port. • <i>active</i> - Logical HA link is active. • <i>reserved</i> - Logical HA link is active, but the physical link is down.

takeover

This represents the state of the node that is taking over storage from its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
giveback	giveback	Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
partners	array[partners]	Nodes in this node's High Availability (HA) group.
ports	array[ports]	
takeover	takeover	This represents the state of the node that is taking over storage from its HA partner.

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

metric

CPU performance for the nodes.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
processor_utilization	integer	Average CPU Utilization for the node

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

metrocluster

Metrocluster

Name	Type	Description
type	string	The Metrocluster configuration type

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to "true" to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

statistics

Raw CPU performance for the nodes.

Name	Type	Description
processor_utilization_base	integer	Base counter for CPU Utilization.
processor_utilization_raw	integer	Raw CPU Utilization for the node. This should be divided by the processor_utilization_base to calculate the percentage CPU utilization for the node.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.

Name	Type	Description
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

vm

Name	Type	Description
provider_type	string	Cloud provider where the VM is hosted.

nodes

Complete node information

Name	Type	Description
_links	_links	
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none"> • example: 2019-04-17T11:49:26-04:00 • format: date-time • readOnly: 1
ha	ha	
location	string	
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
metric	metric	CPU performance for the nodes.
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	

Name	Type	Description
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
statistics	statistics	Raw CPU performance for the nodes.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.
vm	vm	

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

timezone

Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:

- console messages;
- logging to internal ONTAP log files; and
- localized REST API full ISO-8601 date, time, and time zone format information. Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.

Name	Type	Description
name	string	<p>The ONTAP time zone name or identification in either IANA time zone format "Area/Location", or an ONTAP traditional time zone.</p> <p>The initial first node in cluster setting for time zone is "Etc/UTC". "Etc/UTC" is the IANA timezone "Area/Location" specifier for Coordinated Universal Time (UTC), which is an offset of 0.</p> <p>IANA time zone format</p> <p>The IANA time zone, formatted as "Area/Location", is based on geographic areas that have had the same time zone offset for many years.</p> <p>"Location" represents a compound name using additional forward slashes.</p> <p>An example of the "Area/Location" time zone is "America/New_York" and represents most of the United States Eastern Time Zone. Examples of "Area/Location" with "Location" as a compound name are "America/Argentina/Buenos_Aires" and "America/Indiana/Indianapolis".</p> <p>ONTAP traditional time zone</p> <p>Examples of the traditional time zones are "EST5EDT" for the United States Eastern Time Zone and "CET" for Central European Time Zone.</p> <ul style="list-style-type: none"> • example: America/New_York • Introduced in: 9.7

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a cluster configuration

PATCH /cluster

Introduced In: 9.6

Updates the cluster configuration after the cluster is created.

Related ONTAP commands *cluster identity modify*system node modify*vserver services dns modify*vserver services name-service dns modify*timezone*security ssl modify

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
certificate	certificate	Certificate used by cluster and node management interfaces for TLS connection requests.
contact	string	

Name	Type	Description
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
location	string	
management_interfaces	array[management_interfaces]	
name	string	
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
san_optimized	boolean	Specifies if this cluster is an All SAN Array.

Name	Type	Description
timezone	timezone	<p>Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:</p> <ul style="list-style-type: none"> • console messages; • logging to internal ONTAP log files; and • localized REST API full ISO-8601 date, time, and time zone format information. <p>Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.</p> <ul style="list-style-type: none"> • Introduced in: 9.7
uuid	string	

Example request

```
{
  "certificate": {
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "location": "building 1",
  "management_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "name": "cluster1",
  "name_servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "timezone": {
    "name": "America/New_York"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
3604491	Updating timezone failed.
3604520	Internal error. System state is not correct to read or change timezone.
8847361	Too many DNS domains provided.
8847362	Too many name servers provided.
9240587	A name must be provided.
12451843	Certificate does not exist.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

certificate

Certificate used by cluster and node management interfaces for TLS connection requests.

Name	Type	Description
name	string	Certificate name
uuid	string	Certificate UUID

configuration_backup

Name	Type	Description
password	string	
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	
validate_certificate	boolean	Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default.

license

License keys or NLF contents.

Name	Type	Description
keys	array[string]	

ip

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

management_interface

The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	ip	Object to setup an interface along with its default router.

ip

IP information

management_interfaces

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

cluster_interfaces

Network interface

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	List of FRUs on the node. Only returned when requested by name.

Name	Type	Description
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

failure

Indicates the failure code and message.

Name	Type	Description
code	integer	Message code
message	string	Detailed message based on the state.

giveback

Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

partners

Name	Type	Description
name	string	
uuid	string	

ports

Name	Type	Description
number	integer	HA port number

Name	Type	Description
state	string	<p>HA port state:</p> <ul style="list-style-type: none"> • <i>down</i> - Logical HA link is down. • <i>initialized</i> - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port. • <i>armed</i> - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port. • <i>active</i> - Logical HA link is active. • <i>reserved</i> - Logical HA link is active, but the physical link is down.

takeover

This represents the state of the node that is taking over storage from its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
giveback	giveback	Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
partners	array[partners]	Nodes in this node's High Availability (HA) group.
ports	array[ports]	
takeover	takeover	This represents the state of the node that is taking over storage from its HA partner.

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

metric

CPU performance for the nodes.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
processor_utilization	integer	Average CPU Utilization for the node

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

metrocluster

Metrocluster

Name	Type	Description
type	string	The Metrocluster configuration type

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to "true" to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

statistics

Raw CPU performance for the nodes.

Name	Type	Description
processor_utilization_base	integer	Base counter for CPU Utilization.
processor_utilization_raw	integer	Raw CPU Utilization for the node. This should be divided by the processor_utilization_base to calculate the percentage CPU utilization for the node.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.

Name	Type	Description
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

vm

Name	Type	Description
provider_type	string	Cloud provider where the VM is hosted.

nodes

Complete node information

Name	Type	Description
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none"> • example: 2019-04-17T11:49:26-04:00 • format: date-time • readOnly: 1
ha	ha	
location	string	
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	

Name	Type	Description
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.
vm	vm	

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.

Name	Type	Description
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

timezone

Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:

- console messages;
- logging to internal ONTAP log files; and
- localized REST API full ISO-8601 date, time, and time zone format information.
Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.

Name	Type	Description
name	string	<p>The ONTAP time zone name or identification in either IANA time zone format "Area/Location", or an ONTAP traditional time zone.</p> <p>The initial first node in cluster setting for time zone is "Etc/UTC". "Etc/UTC" is the IANA timezone "Area/Location" specifier for Coordinated Universal Time (UTC), which is an offset of 0.</p> <p>IANA time zone format</p> <p>The IANA time zone, formatted as "Area/Location", is based on geographic areas that have had the same time zone offset for many years.</p> <p>"Location" represents a compound name using additional forward slashes.</p> <p>An example of the "Area/Location" time zone is "America/New_York" and represents most of the United States Eastern Time Zone. Examples of "Area/Location" with "Location" as a compound name are "America/Argentina/Buenos_Aires" and "America/Indiana/Indianapolis".</p> <p>ONTAP traditional time zone</p> <p>Examples of the traditional time zones are "EST5EDT" for the United States Eastern Time Zone and "CET" for Central European Time Zone.</p> <ul style="list-style-type: none"> • example: America/New_York • Introduced in: 9.7

cluster

Complete cluster information

Name	Type	Description
certificate	certificate	Certificate used by cluster and node management interfaces for TLS connection requests.
contact	string	
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
location	string	
management_interfaces	array[management_interfaces]	
name	string	
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
san_optimized	boolean	Specifies if this cluster is an All SAN Array.

Name	Type	Description
timezone	timezone	<p>Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:</p> <ul style="list-style-type: none"> • console messages; • logging to internal ONTAP log files; and • localized REST API full ISO-8601 date, time, and time zone format information. Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC. <ul style="list-style-type: none"> • Introduced in: 9.7
uuid	string	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a cluster

POST /cluster

Introduced In: 9.6

Creates a cluster.

Required properties

- name
- password

Recommended optional properties

- location
- contact
- dns_domains
- name_servers
- ntp_servers
- license
- configuration_backup
- management_interface
- nodes
- timezone

Learn more

- [DOC /cluster](#)

Parameters

Name	Type	In	Required	Description
single_node_cluster	boolean	query	False	Configures a single node cluster. All cluster ports are reassigned to the default network. The storage failover settings are configured to non-HA. The node reboots during this operation.
create_recommended_aggregates	boolean	query	False	<p>Create aggregates based on an optimal layout recommended by the system.</p> <ul style="list-style-type: none"> • Introduced in: 9.7 • Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
configuration_backup	configuration_backup	
contact	string	

Name	Type	Description
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
license	license	License keys or NLF contents.
location	string	
management_interface	management_interface	The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[management_interfaces]	
name	string	
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
nodes	array[nodes]	

Name	Type	Description
ntp_servers	array[string]	Host name, IPv4 address, or IPv6 address for the external NTP time servers.
password	string	Initial admin password used to create the cluster.
san_optimized	boolean	Specifies if this cluster is an All SAN Array.
timezone	timezone	<p>Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:</p> <ul style="list-style-type: none"> • console messages; • logging to internal ONTAP log files; and • localized REST API full ISO-8601 date, time, and time zone format information. <p>Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.</p> <ul style="list-style-type: none"> • Introduced in: 9.7
uuid	string	

Example request

```
{
  "configuration_backup": {
    "password": "yourpassword",
    "url": "http://10.224.65.198/backups",
    "username": "me"
  },
  "contact": "<a href="
mailto:support@company.com">support@company.com</a>",
  "dns_domains": [
    "example.com",
    "example2.example3.com"
  ],
  "license": {
    "keys": [
      "AMEPOSIOIKLKGEEEEEDGNDEKSJDE"
    ]
  },
  "location": "building 1",
  "management_interface": {
    "ip": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    }
  },
  "management_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "name": "cluster1",
  "name_servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "nodes": [
    {
      "cluster_interface": {
        "ip": {
          "address": "10.10.10.7"
        }
      },
      "cluster_interfaces": [
```

```

    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "controller": {
    "flash_cache": [
      {
        "capacity": "1024000000000",
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
        "slot": "6-1",
        "state": "string"
      }
    ],
    "frus": [
      {
        "id": 0,
        "state": "string",
        "type": "string"
      }
    ],
    "over_temperature": "string"
  },
  "date": "2019-04-17T11:49:26-04:00",
  "ha": {
    "giveback": {
      "failure": {
        "code": "852126",
        "message": "Failed to initiate giveback. Run the \"storage failover show-giveback\" command for more information."
      },
      "state": "failed"
    },
    "partners": [
      {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "ports": [
      {
        "number": "0",

```

```

        "state": "active"
    }
],
"takeover": {
    "failure": {
        "code": "852130",
        "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
    },
    "state": "failed"
}
},
"location": "rack 2 row 5",
"management_interface": {
    "ip": {
        "address": "10.10.10.7"
    }
},
"management_interfaces": [
    {
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
    }
],
"membership": "string",
"metrocluster": {
    "type": "string"
},
"model": "FAS3070",
"name": "node-01",
"serial_number": "4048820-60-9",
"service_processor": {
    "dhcp_enabled": null,
    "firmware_version": "string",
    "ipv4_interface": {
        "address": "10.10.10.7",
        "gateway": "10.1.1.1",
        "netmask": "24"
    },
    "link_status": "string",
    "mac_address": "string",
    "state": "string"
},
"state": "string",
"system_id": "0537035403",
"system_machine_type": "7Y56-CTOWW1",

```

```

    "uptime": "300536",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
    "vendor_serial_number": "791603000068",
    "vm": {
      "provider_type": "string"
    }
  },
  "ntp_servers": [
    "time.nist.gov",
    "10.98.19.20",
    "2610:20:6F15:15::27"
  ],
  "password": "mypassword",
  "timezone": {
    "name": "America/New_York"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

{
  "job": {
    "uuid": "string"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
262245	The value provided is invalid.
1179813	Fields set for one node must be set for all nodes.
1179817	The IP address, subnet mask, and gateway must all be provided for cluster management interface.
1179818	The IP address and gateway must be of the same family.
1179821	An IP address and subnet mask conflicts with an existing entry.
1179824	An invalid gateway was provided.
1179825	All management and cluster config IP addresses must belong to the same address family.
2097165	An NTP server could not be reached.
8847361	Too many DNS domains provided.
8847362	Too many name servers provided.
8847394	An invalid DNS domain was provided.
8978433	An invalid license key was provided.
9240587	A name must be provided.
9240594	An invalid name was provided.
39387137	The URL provided is invalid.
131727360	A node could not be added to the cluster. This is a generic code, see response message for details.
131727388	Hostnames for NTP servers cannot be used without DNS configured.
131727389	URL and username are required for configuration backup.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

certificate

Certificate used by cluster and node management interfaces for TLS connection requests.

Name	Type	Description
name	string	Certificate name
uuid	string	Certificate UUID

configuration_backup

Name	Type	Description
password	string	
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	
validate_certificate	boolean	Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default.

license

License keys or NLF contents.

Name	Type	Description
keys	array[string]	

ip

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

management_interface

The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration.

Name	Type	Description
ip	ip	Object to setup an interface along with its default router.

ip

IP information

management_interfaces

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

cluster_interfaces

Network interface

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.

Name	Type	Description
frus	array[frus]	List of FRUs on the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

failure

Indicates the failure code and message.

Name	Type	Description
code	integer	Message code
message	string	Detailed message based on the state.

giveback

Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

partners

Name	Type	Description
name	string	
uuid	string	

ports

Name	Type	Description
number	integer	HA port number

Name	Type	Description
state	string	<p>HA port state:</p> <ul style="list-style-type: none"> • <i>down</i> - Logical HA link is down. • <i>initialized</i> - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port. • <i>armed</i> - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port. • <i>active</i> - Logical HA link is active. • <i>reserved</i> - Logical HA link is active, but the physical link is down.

takeover

This represents the state of the node that is taking over storage from its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
giveback	giveback	Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
partners	array[partners]	Nodes in this node's High Availability (HA) group.
ports	array[ports]	
takeover	takeover	This represents the state of the node that is taking over storage from its HA partner.

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

metric

CPU performance for the nodes.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
processor_utilization	integer	Average CPU Utilization for the node

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

metrocluster

Metrocluster

Name	Type	Description
type	string	The Metrocluster configuration type

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

statistics

Raw CPU performance for the nodes.

Name	Type	Description
processor_utilization_base	integer	Base counter for CPU Utilization.

Name	Type	Description
processor_utilization_raw	integer	Raw CPU Utilization for the node. This should be divided by the processor_utilization_base to calculate the percentage CPU utilization for the node.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.

Name	Type	Description
minor	integer	The minor portion of the version.

vm

Name	Type	Description
provider_type	string	Cloud provider where the VM is hosted.

nodes

Complete node information

Name	Type	Description
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none"> • example: 2019-04-17T11:49:26-04:00 • format: date-time • readOnly: 1
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	

Name	Type	Description
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.
vm	vm	

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.

Name	Type	Description
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

timezone

Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:

- console messages;
- logging to internal ONTAP log files; and
- localized REST API full ISO-8601 date, time, and time zone format information.
Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC.

Name	Type	Description
name	string	<p>The ONTAP time zone name or identification in either IANA time zone format "Area/Location", or an ONTAP traditional time zone.</p> <p>The initial first node in cluster setting for time zone is "Etc/UTC". "Etc/UTC" is the IANA timezone "Area/Location" specifier for Coordinated Universal Time (UTC), which is an offset of 0.</p> <p>IANA time zone format</p> <p>The IANA time zone, formatted as "Area/Location", is based on geographic areas that have had the same time zone offset for many years.</p> <p>"Location" represents a compound name using additional forward slashes.</p> <p>An example of the "Area/Location" time zone is "America/New_York" and represents most of the United States Eastern Time Zone. Examples of "Area/Location" with "Location" as a compound name are "America/Argentina/Buenos_Aires" and "America/Indiana/Indianapolis".</p> <p>ONTAP traditional time zone</p> <p>Examples of the traditional time zones are "EST5EDT" for the United States Eastern Time Zone and "CET" for Central European Time Zone.</p> <ul style="list-style-type: none"> • example: America/New_York • Introduced in: 9.7

cluster

Complete cluster information

Name	Type	Description
configuration_backup	configuration_backup	
contact	string	
dns_domains	array[string]	<p>A list of DNS domains. Domain names have the following requirements:</p> <ul style="list-style-type: none"> • The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_". • The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9. • The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9. • The top level domain must contain only the following characters: A through Z, a through z. • The system reserves the following names: "all", "local", and "localhost".
license	license	License keys or NLF contents.
location	string	
management_interface	management_interface	The management interface of the cluster. The subnet mask and gateway for this interface are used for the node management interfaces provided in the node configuration.
management_interfaces	array[management_interfaces]	
name	string	

Name	Type	Description
name_servers	array[string]	The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
nodes	array[nodes]	
ntp_servers	array[string]	Host name, IPv4 address, or IPv6 address for the external NTP time servers.
password	string	Initial admin password used to create the cluster.
san_optimized	boolean	Specifies if this cluster is an All SAN Array.
timezone	timezone	<p>Provides the cluster-wide time zone information that localizes time found on messages displayed on each node's:</p> <ul style="list-style-type: none"> • console messages; • logging to internal ONTAP log files; and • localized REST API full ISO-8601 date, time, and time zone format information. Machine-to-machine interfaces, such as file access protocols (NFS, CIFS), block access protocols (SAN), and other protocols such as Manage ONTAP (ONTAPI), use second or subsecond time values that are based on world time or UTC. <ul style="list-style-type: none"> • Introduced in: 9.7
uuid	string	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve cluster chassis

Cluster chassis endpoint overview

Overview

You can use the chassis GET API to retrieve all of the chassis information in the cluster.

Examples

Retrieving a list of chassis from the cluster

The following example shows the response with a list of chassis in the cluster:

```
# The API:
/api/cluster/chassis

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/chassis" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "id": "021352005981",
      "_links": {
        "self": {
          "href": "/api/cluster/chassis/021352005981"
        }
      }
    },
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/cluster/chassis"
    }
  }
}
```

Retrieving a specific chassis from the cluster

The following example shows the response of the requested chassis. If there is no chassis with the requested ID, an error is returned.

```
# The API:
/api/cluster/chassis/{id}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/chassis/021352005981" -H
"accept: application/hal+json"

# The response:
{
  "id": "021352005981",
  "state": "ok",
```



```

"nodes": [
  {
    "name": "node-1",
    "uuid": "6ede364b-c3d0-11e8-a86a-00a098567f31",
    "position": "top",
    "usbs": {
      "supported": true,
      "enabled": true,
      "ports": [
        {
          "connected": false
        }
      ]
    },
    "pcis": {
      "cards": [
        {
          "slot": "0",
          "device": "Gigabit Ethernet I210",
          "info": "\t e0M MAC Address:    d0:39:ea:3f:06:2b (auto-1000t-
fd-up) \n\t e0S MAC Address:    d0:39:ea:3f:06:2c (auto-1000t-fd-up) \n\t
Device Type:          1533\n\t Firmware Version:    3.25-0.0 0x800005D1\n"
        },
        {
          "slot": "0",
          "device": "Intel Lewisburg series chipset SATA Controller",
          "info": "\t Additional Info: 0 (0xaaf00000)  \n\t
SHM2S86Q120GLM22NP FW1146 114473MB 512B/sect (SPG190108HJ)  \n"
        }
      ]
    },
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/6ede364b-c3d0-11e8-a86a-00a098567f31"
      }
    }
  },
  {
    "frus": [
      {
        "id": "PSU2",
        "type": "psu",
        "state": "ok"
      },
      {
        "id": "PSU1",

```

```

    "type": "psu",
    "state": "ok"
  },
  {
    "id": "Fan2",
    "type": "fan",
    "state": "ok"
  },
  {
    "id": "Fan3",
    "type": "fan",
    "state": "ok"
  },
  {
    "id": "Fan1",
    "type": "fan",
    "state": "ok"
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/chassis/021352005981"
  }
}
}

```

Retrieve a collection of chassis

GET /cluster/chassis

Introduced In: 9.6

Retrieves a collection of chassis.

Related ONTAP commands

- `system chassis show`
- `system chassis fru show`

Learn more

- [DOC /cluster/chassis](#)

Parameters

Name	Type	In	Required	Description
frus.state	string	query	False	Filter by frus.state

Name	Type	In	Required	Description
frus.id	string	query	False	Filter by frus.id
frus.type	string	query	False	Filter by frus.type
shelves.uid	string	query	False	Filter by shelves.uid
nodes.name	string	query	False	Filter by nodes.name
nodes.uuid	string	query	False	Filter by nodes.uuid
nodes.position	string	query	False	Filter by nodes.position <ul style="list-style-type: none"> Introduced in: 9.8
state	string	query	False	Filter by state
id	string	query	False	Filter by id
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records.
records	array[chassis]	

Example response

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "frus": [
        {
          "id": "string",
          "state": "string",
          "type": "string"
        }
      ],
      "id": "021352005981",
      "nodes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "position": "top",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "shelves": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "uid": "7777841915827391056"
        }
      ],
      "state": "string"
    }
  ]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

frus

Name	Type	Description
id	string	
state	string	
type	string	

_links

Name	Type	Description
self	href	

nodes

Name	Type	Description
_links	_links	
name	string	
position	string	The position of the node in the chassis.
uuid	string	

shelf_reference

Shelf

Name	Type	Description
_links	_links	
uid	string	

chassis

Name	Type	Description
frus	array[frus]	List of FRUs in the chassis.
id	string	
nodes	array[nodes]	List of nodes in the chassis.
shelves	array[shelf_reference]	List of shelves in chassis.
state	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a chassis

GET /cluster/chassis/{id}

Introduced In: 9.6

Retrieves a specific chassis.

Related ONTAP commands

- `system chassis show`
- `system chassis fru show`

Learn more

- [DOC /cluster/chassis](#)

Parameters

Name	Type	In	Required	Description
id	string	path	True	Chassis ID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
frus	array[frus]	List of FRUs in the chassis.
id	string	
nodes	array[nodes]	List of nodes in the chassis.
shelves	array[shelf_reference]	List of shelves in chassis.
state	string	

Example response

```
{
  "frus": [
    {
      "id": "string",
      "state": "string",
      "type": "string"
    }
  ],
  "id": "021352005981",
  "nodes": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "position": "top",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "shelves": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uid": "7777841915827391056"
    }
  ],
  "state": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

frus

Name	Type	Description
id	string	
state	string	
type	string	

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

nodes

Name	Type	Description
_links	_links	
name	string	
position	string	The position of the node in the chassis.
uuid	string	

shelf_reference

Shelf

Name	Type	Description
_links	_links	
uid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View and manage cluster jobs

Cluster jobs endpoint overview

Overview

You can use this API to view and manipulate jobs. Jobs provide information about asynchronous operations. Some long-running jobs are paused or cancelled by calling a PATCH request. Individual operations indicate if they support PATCH requests on the job. After a job transitions to a terminal state, it is deleted after a default time of 300 seconds. Attempts to call a GET or PATCH request on the job returns a 404 error code After the job has been deleted.

Example

The following examples show how to retrieve and update a job state:

Retrieving job information

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/b5145e1d-b53b-11e8-8252-005056bbd8f5" -H "accept: application/json"

# The response:
{
  "uuid": "b5145e1d-b53b-11e8-8252-005056bbd8f5",
  "code": 0,
  "description": "Cluster Backup Job",
  "state": "running",
  "message": "creating_node_backups",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/b5145e1d-b53b-11e8-8252-005056bbd8f5"
    }
  }
}
```

Updating a job that supports the new state

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/cluster/jobs/b5145e1d-b53b-11e8-8252-005056bbd8f5?action=cancel" -H "accept: application/json"
```

Retrieve recent asynchronous jobs

GET /cluster/jobs

Introduced In: 9.6

Retrieves a list of recently running asynchronous jobs. After a job transitions to a failure or success state, it is deleted after a default time of 300 seconds.

Parameters

Name	Type	In	Required	Description
message	string	query	False	Filter by message
uuid	string	query	False	Filter by uuid
code	integer	query	False	Filter by code
description	string	query	False	Filter by description
end_time	string	query	False	Filter by end_time
svm.uuid	string	query	False	Filter by svm.uuid • Introduced in: 9.8
svm.name	string	query	False	Filter by svm.name • Introduced in: 9.8
start_time	string	query	False	Filter by start_time
state	string	query	False	Filter by state
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	
records	array[job]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "1",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "code": "0",
      "description": "App Snapshot Job",
      "end_time": "string",
      "message": "Complete: Successful",
      "start_time": "string",
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

job

Name	Type	Description
_links	_links	
code	integer	If the state indicates "failure", this is the final error code.
description	string	The description of the job to help identify it independent of the UUID.
end_time	string	The time the job ended.
message	string	A message corresponding to the state of the job providing additional details about the current state.

Name	Type	Description
start_time	string	The time the job started.
state	string	The state of the job.
svm	svm	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve details of an asynchronous job

GET `/cluster/jobs/{uuid}`

Introduced In: 9.6

Retrieves the details of a specific asynchronous job. After a job transitions to a failure or success state, it is deleted after a default time of 300 seconds.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Job UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
code	integer	If the state indicates "failure", this is the final error code.
description	string	The description of the job to help identify it independent of the UUID.
end_time	string	The time the job ended.
message	string	A message corresponding to the state of the job providing additional details about the current state.
start_time	string	The time the job started.
state	string	The state of the job.
svm	svm	
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "code": "0",
  "description": "App Snapshot Job",
  "end_time": "string",
  "message": "Complete: Successful",
  "start_time": "string",
  "state": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the state of an asynchronous job

PATCH /cluster/jobs/{uuid}

Introduced In: 9.6

Updates the state of a specific asynchronous job.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Job UUID
action	string	query	False	<p>Requests a job to pause, resume, or cancel. Note that not all jobs support these actions. A job can only be resumed if it is in a paused state. After you successfully request a job to be cancelled, the job state changes to either success or failure.</p> <ul style="list-style-type: none">enum: ["pause", "resume", "cancel"]

Response

```
Status: 200, Ok
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
459753	Command execution failed with custom error from the program.
458762	Job is already in a terminal state.
458773	The Job Manager is not initialized.
458771	The specified job is running.
458776	The specified job is not currently running.

Error Code	Description
458783	This job does not support pause.
458784	This job does not support cancel.

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve capacity pool licenses

Cluster licensing capacity-pools endpoint overview

Overview

Capacity pool licenses are installed on and managed by the license manager. Each ONTAP node that is using the capacity pools licensing model is associated with a capacity pool license from which capacity is leased for data aggregates.

This API is used to retrieve information about associations between ONTAP nodes in the cluster and capacity pool licenses. It also reports how much capacity each node is consuming from the capacity pool.

Examples

Retrieving a collection of capacity pools associated with the cluster

This example retrieves a collection that contains two capacity pool licenses, each of which is associated with an HA pair of nodes in a four-node cluster.

```
# API
curl -X GET "https://<mgmt-ip>/api/cluster/licensing/capacity-pools"

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "serial_number": "390000100",
      "license_manager": {
        "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566",
        "_links": {
          "self": {
            "href": "/api/cluster/licensing/license-managers/4ea7a442-86d1-11e0-ae1c-112233445566"
          }
        }
      }
    },
    {
      "nodes": [
        {
          "node": {
            "name": "node-1",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563411"
          },
          "used_size": 109951162776,
          "_links": {
            "self": {
              "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-123478563411"
            }
          }
        },
        {
          "node": {
            "name": "node-2",
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          },
          "used_size": 109951162776,
          "_links": {
            "self": {
```

```

    "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-123478563412"
  }
}
],
"_links":{
  "self":{
    "href":"/api/cluster/licensing/capacity-pools/390000100"
  }
}
},
{
  "serial_number":"390000101",
  "license_manager": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566",
    "_links": {
      "self": {
        "href": "/api/cluster/licensing/license-managers/4ea7a442-86d1-11e0-ae1c-112233445566"
      }
    }
  },
  "nodes":[
    {
      "node":{
        "name":"node-3",
        "uuid":"4ea7a442-86d1-11e0-ae1c-123478563413"
      },
      "used_size":219902325552,
      "_links":{
        "self":{
          "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-123478563413"
        }
      }
    },
    {
      "node":{
        "name":"node-4",
        "uuid":"4ea7a442-86d1-11e0-ae1c-123478563414"
      },
      "used_size":219902325552,
      "_links":{
        "self":{
          "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-

```

```

123478563414"
    }
  }
},
"_links":{
  "self":{
    "href":"/api/cluster/licensing/capacity-pools/390000101"
  }
}
],
"num_records":2,
"_links":{
  "self":{
    "href":"/api/cluster/licensing/capacity-pools"
  }
}
}
}

```

Retrieving information about nodes associated with a specific capacity pool license

This example retrieves information about the nodes that are associated with a capacity pool license of the serial number 390000100.

```

# API
curl -X GET "https://<mgmt-ip>/api/cluster/licensing/capacity-
pools/390000100"

# Response
200 OK

# JSON Body
{
  "serial_number":"390000100",
  "license_manager": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566",
    "_links": {
      "self": {
        "href": "/api/cluster/licensing/license-managers/4ea7a442-86d1-11e0-
ae1c-112233445566"
      }
    }
  },
  "nodes": [
    {

```

```

    "node":{
      "name":"node-1",
      "uuid":"4ea7a442-86d1-11e0-ae1c-123478563411"
    },
    "used_size":1099511627776,
    "_links":{
      "self":{
        "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-123478563411"
      }
    }
  },
  {
    "node":{
      "name":"node-2",
      "uuid":"4ea7a442-86d1-11e0-ae1c-123478563412"
    },
    "used_size":1099511627776,
    "_links":{
      "self":{
        "href": "/api/cluster/nodes/4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
  }
],
"_links":{
  "self":{
    "href":"/api/cluster/licensing/capacity-pools/390000100"
  }
}
}

```

Retrieve capacity pools

GET /cluster/licensing/capacity-pools

Introduced In: 9.8

Retrieves a collection of capacity pools.

Learn more

- [DOC /cluster/licensing/capacity-pools](#)

Related ONTAP commands

- system license show-status
- system license show

Parameters

Name	Type	In	Required	Description
nodes.used_size	integer	query	False	Filter by nodes.used_size
nodes.node.name	string	query	False	Filter by nodes.node.name
nodes.node.uuid	string	query	False	Filter by nodes.node.uuid
license_manager.uuid	string	query	False	Filter by license_manager.uuid
serial_number	string	query	False	Filter by serial_number
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	collection_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "license_manager": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
      },
      "nodes": [
        {
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "used_size": 0
        }
      ],
      "serial_number": "390000100"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

collection_links

Name	Type	Description
next	href	
self	href	

self_link

Name	Type	Description
self	href	

license_manager

License manager instance where this capacity pool license in installed.

Name	Type	Description
_links	self_link	
uuid	string	

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

nodes

Information on a node from the capacity licensing perspective.

Name	Type	Description
node	node_reference	

Name	Type	Description
used_size	integer	Capacity, in bytes, that is currently used by the node.

records

Information on a capacity pool license and how it is associated with the cluster.

Name	Type	Description
_links	self_link	
license_manager	license_manager	License manager instance where this capacity pool license is installed.
nodes	array[nodes]	Nodes in the cluster associated with this capacity pool.
serial_number	string	Serial number of the capacity pool license.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve capacity pool information

GET /cluster/licensing/capacity-pools/{serial_number}

Introduced In: 9.8

Retrieves information about the capacity pool.

Learn more

- [DOC /cluster/licensing/capacity-pools](#)

Related ONTAP commands

- `system license show-status`
- `system license show`

Parameters

Name	Type	In	Required	Description
serial_number	string	path	True	Serial number of the capacity pool license.
nodes.used_size	integer	query	False	Filter by nodes.used_size
nodes.node.name	string	query	False	Filter by nodes.node.name
nodes.node.uuid	string	query	False	Filter by nodes.node.uuid
license_manager.uuid	string	query	False	Filter by license_manager.uuid
serial_number	string	query	False	Filter by serial_number
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	self_link	

Name	Type	Description
license_manager	license_manager	License manager instance where this capacity pool license is installed.
nodes	array[nodes]	Nodes in the cluster associated with this capacity pool.
serial_number	string	Serial number of the capacity pool license.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "license_manager": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
  },
  "nodes": [
    {
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "used_size": 0
    }
  ],
  "serial_number": "390000100"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

license_manager

License manager instance where this capacity pool license in installed.

Name	Type	Description
_links	self_link	
uuid	string	

_links

Name	Type	Description
self	href	

node_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

nodes

Information on a node from the capacity licensing perspective.

Name	Type	Description
node	node_reference	
used_size	integer	Capacity, in bytes, that is currently used by the node.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster license managers

Cluster licensing license-managers endpoint overview

Overview

This API is used to manage information about the license manager instance associated with the cluster.

When an ONTAP cluster is initially created to use the capacity pools licensing model, information about the license manager instance that the cluster should use is pre-configured. Generally, this configuration does not need to be updated unless the license manager instance changes its IP address.

The license manager is currently bundled with the ONTAP Select Deploy utility and runs on the same VM as ONTAP Select Deploy. Use this API to update the license manager IP address when the Deploy VM changes its IP address.

Examples

Retrieving information about the license manager instance associated with the cluster

```
# API
curl -X GET "https://<mgmt-ip>/api/cluster/licensing/license-managers"

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566",
      "uri": {
        "host": "10.1.1.1",
      },
      "default": true
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/cluster/licensing/license-managers"
    }
  }
}
```

Updating an existing license manager instance

```
# API
curl -X PATCH "https://<mgmt-ip>/api/cluster/licensing/license-
managers/4ea7a442-86d1-11e0-ae1c-112233445566"

# JSON Body
{
  "uri": {
    "host": "10.1.1.3"
  }
}

# Response
202 Accepted
```

Retrieve license managers

GET /cluster/licensing/license-managers

Introduced In: 9.8

Retrieves a collection of license managers.

Learn more

- [DOC /cluster/licensing/license-managers](#)

Related ONTAP commands

- `system license license-manager show`

Parameters

Name	Type	In	Required	Description
default	boolean	query	False	Filter by default
uri.host	string	query	False	Filter by uri.host
uuid	string	query	False	Filter by uuid
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	collection_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "uri": {
        "host": "10.1.1.1"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

collection_links

Name	Type	Description
next	href	
self	href	

self_link

Name	Type	Description
self	href	

uri

License manager URI.

Name	Type	Description
host	string	License manager host name, IPv4 or IPv6 address.

records

Information on a license manager instance associated with the cluster.

Name	Type	Description
_links	self_link	
default	boolean	Flag that indicates whether it's the default license manager instance used by the cluster.' When a capacity pool is created and if the license manager field is omitted, it is assumed that the license of the capacity pool is installed on the default license manager instance.
uri	uri	License manager URI.
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve license manager information

GET /cluster/licensing/license-managers/{uuid}

Introduced In: 9.8

Retrieves information about the license manager.

Learn more

- [DOC /cluster/licensing/license-managers](#)

Related ONTAP commands

- `system license license-manager show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	
default	boolean	query	False	Filter by default
uri.host	string	query	False	Filter by uri.host
uuid	string	query	False	Filter by uuid

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	self_link	
default	boolean	Flag that indicates whether it's the default license manager instance used by the cluster. When a capacity pool is created and if the license manager field is omitted, it is assumed that the license of the capacity pool is installed on the default license manager instance.
uri	uri	License manager URI.
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "uri": {
    "host": "10.1.1.1"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

uri

License manager URI.

Name	Type	Description
host	string	License manager host name, IPv4 or IPv6 address.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the license manager configuration

PATCH /cluster/licensing/license-managers/{uuid}

Introduced In: 9.8

Updates the license manager configuration.

Learn more

- [DOC /cluster/licensing/license-managers](#)

Related ONTAP commands

- `system license license-manager modify`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Request Body

Name	Type	Description
default	boolean	Flag that indicates whether it's the default license manager instance used by the cluster. When a capacity pool is created and if the license manager field is omitted, it is assumed that the license of the capacity pool is installed on the default license manager instance.
uri	uri	License manager URI.
uuid	string	

Example request

```
{
  "uri": {
    "host": "10.1.1.1"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-112233445566"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1115532	The requested update to the license manager information failed.

Name	Type	Description
errors	array[error]	

Example error

```
{
  "errors": [
    {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  ]
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

uri

License manager URI.

Name	Type	Description
host	string	License manager host name, IPv4 or IPv6 address.

license_manager

Information on a license manager instance associated with the cluster.

Name	Type	Description
default	boolean	Flag that indicates whether it's the default license manager instance used by the cluster.' When a capacity pool is created and if the license manager field is omitted, it is assumed that the license of the capacity pool is installed on the default license manager instance.
uri	uri	License manager URI.
uuid	string	

_links

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster licensing

Cluster licensing licenses endpoint overview

Overview

Licensing allows you to tailor a system to meet an organization's specific needs. You can enable new features by purchasing a license from a NetApp sales associate. After installation of the license, the new feature is available immediately.

This interface manages licenses according to their supported feature. By default, the interface displays packages with installed licenses, but you can also return unlicensed packages.

Each feature has a compliance state that is indicated at the package level. Individual licenses also contain a compliance state indicated in the "licenses" array. The state of the package is determined by analyzing the underlying licenses according to the following criteria:

- Licensing terms
- Cluster state

Licensing terms

The licensing terms define the conditions under which a package is considered "compliant". Individual licenses are evaluated based on the following:

- Scope
- Time period
- Usage

Scope

A package can be licensed under the following scopes:

- Site - Permits the feature to be used by any node that is a member of the cluster.
- Cluster - Permits the feature to be used by any node that is a member of the cluster.
- Node - Permits the authorized node to use the feature. Within a cluster, if you don't supply every node with a valid license, the package state indicates "noncompliant". You must purchase a license for each node in a cluster for the package to be considered "compliant".

Time period

Some package licenses are only valid for a limited period of time. After a license has expired, the package state changes to "noncompliant". You need to purchase a new license for the package to return to a "compliant" state.

Usage

Some package licenses have additional terms that need to be maintained to keep a license in compliance. These conditions are defined by the individual license. For example, a license might define the maximum amount of storage that a node can allocate for the license to be "compliant".

Cluster state

A cluster's state consists of the following:

- Node online status
- Node cluster membership

Some features require that a node be online to display a valid compliance state. If a node cannot be reached or is not known to the cluster, the individual license might indicate an "unknown" state.

Licensing keys

A license is issued in one of the following two formats:

- 28-character key
- NetApp License File (NLF)

The following is an example of a 28-character key:

```
AMEPOSIOIKLKGEEEEEDGNDEKSJDEEE
```

The following is an example of an NLF key:

```
{
  "statusResp": {
    "version": "1",
    "serialNumber": "123456789",
    "message": "Success",
    "licenses": {
      "capacity": "1",
      "type": "capacity",
      "licenseProtocol": "FABRICPOOL-TB",
      "package": "FabricPool",
      "licenseScope": "cluster"
    },
    "snStatus": "Active",
    "product": "fabricpool",
    "statusCode": "S007"
  },
  "Signature": "signatureABC"
}
```

You can use this API to submit either format to enable features.

Examples

Retrieving a collection of licenses organized by package

This example retrieves a collection that contains one entry for each package (filtered to only the 'fabricpool' package).

```

# API
curl -X GET "https://<mgmt-
ip>/api/cluster/licensing/licenses/?fields=*&name=fabricpool"

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "name": "fabricpool",
      "scope": "cluster",
      "state": "compliant",
      "licenses": [
        {
          "owner": "testcluster-1",
          "serial_number": "4149027342",
          "state": "compliant",
          "capacity": {
            "maximum_size": 1099511627776,
            "used_size": 0
          }
        }
      ],
      "_links": {
        "self": {
          "href": "/api/cluster/licensing/licenses/fabricpool"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/cluster/licensing/licenses/?fields=*&name=fabricpool"
    }
  }
}

```

Retrieving a collection of installed licenses

This example retrieves a collection containing all packages (except base) that have installed licenses.

```
# API
```

```
curl -X GET "https://<mgmt-  
ip>/api/cluster/licensing/licenses/?fields=*&name=!base"
```

```
# Response
```

```
200 OK
```

```
# JSON Body
```

```
{  
  "records": [  
    {  
      "name": "nfs",  
      "scope": "node",  
      "state": "compliant",  
      "licenses": [  
        {  
          "owner": "testcluster-1",  
          "serial_number": "1-81-0000000000000004149027492",  
          "state": "compliant"  
        }  
      ],  
      "_links": {  
        "self": {  
          "href": "/api/cluster/licensing/licenses/nfs"  
        }  
      }  
    },  
    {  
      "name": "cifs",  
      "scope": "node",  
      "state": "compliant",  
      "licenses": [  
        {  
          "owner": "testcluster-1",  
          "serial_number": "1-81-0000000000000004149027492",  
          "state": "compliant"  
        }  
      ],  
      "_links": {  
        "self": {  
          "href": "/api/cluster/licensing/licenses/cifs"  
        }  
      }  
    }  
  ],  
  "num_records": 2,  
  "_links": {
```

```
"self": {
  "href": "/api/cluster/licensing/licenses/?fields=*&name=!base"
}
}
}
```

Retrieving a collection of unlicensed packages

By default, unlicensed packages are filtered from the collection output. This example shows how to use a query to retrieve unlicensed packages.

```
# API
curl -X GET "https://<mgmt-
ip>/api/cluster/licensing/licenses?name=flexcache&state=unlicensed"

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "name": "flexcache",
      "_links": {
        "self": {
          "href": "/api/cluster/licensing/licenses/flexcache"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href":
"/api/cluster/licensing/licenses?name=flexcache&state=unlicensed"
    }
  }
}
```

Installing an NLF license

This example installs a single license in the NLF format.



You must escape all the double quotes and backslash characters of the JSON license before it can be placed in the POST request.

```
# API
curl -X POST "https://<mgmt-ip>/api/cluster/licensing/licenses/"

# JSON Body
{
  "keys" : [ "{\"statusResp\":{\"snStatus\": \"Active\", \"licenses\": {
    \"package\": \"FabricPool\", \"capacity\": \"1\", \"licenseProtocol\":
    \"FABRICPOOL-TB\", \"type\": \"capacity\", \"licenseScope\": \"cluster\"},
    \"message\": \"Success\", \"statusCode\": \"S007\", \"version\": \"1\",
    \"product\": \"fabricpool\", \"serialNumber\": \"4149027342\"},
    \"Signature\": \"SignatureABC\"}" ]
}

# Response
201 Created
```

Installing a 28-character key

This example installs a single 28-character key formatted license.

```
# API
curl -X POST "https://<mgmt-ip>/api/cluster/licensing/licenses/"

# JSON Body
{
  "keys" : [ "AAAAAAAAAAAAAAAAAAAAAAAAAAAA" ]
}

# Response
201 Created
```

Installing multiple licenses with one API call

This example shows how multiple keys can be provided to install multiple features in a single API call.

```
# API
curl -X POST "https://<mgmt-ip>/api/cluster/licensing/licenses/"

# JSON Body
{
  "keys" : [ "AAAAAAAAAAAAAAAAAAAAAAAAAAAA",
             "BBBBBBBBBBBBBBBBBBBBBBBBBBBB" ]
}

# Response
201 Created
```

Retrieving information for a specific license package

This example shows how to retrieve information about the specific feature package `fabricpool`.

```
# API
curl -X GET "https://<mgmt-ip>/api/cluster/licensing/licenses/fabricpool/"

# Response
200 OK

# JSON Body
{
  "name": "fabricpool",
  "scope": "cluster",
  "state": "compliant",
  "licenses": [
    {
      "owner": "testcluster-1",
      "serial_number": "123456789",
      "state": "compliant",
      "capacity": {
        "maximum_size": 109951162777600,
        "used_size": 0
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster/licensing/licenses/fabricpool/"
    }
  }
}
```


Deleting a specific license

This example show how to delete a CIFS site license.

```
# API
curl -X DELETE "https://<mgmt-
ip>/api/cluster/licensing/licenses/cifs/?serial_number=1-80-000011"

# JSON Body
{}

# Response
200 OK
```

Deleting with a query

The following example shows how to delete all NFS licenses specified with the '*' query.

```
# API
curl -X DELETE "https://<mgmt-
ip>/api/cluster/licensing/licenses/nfs/?serial_number=*"

# JSON Body
{}

# Response
200 OK
```

Retrieve license packages

GET /cluster/licensing/licenses

Introduced In: 9.6

Retrieves a collection of license packages.



By default, the GET method only returns licensed packages. You must provide the following query "state=unlicensed" to retrieve unlicensed packages.

Related ONTAP commands

- `system license show-status`
- `system license show`

Parameters

Name	Type	In	Required	Description
state	string	query	False	Filter by state
scope	string	query	False	Filter by scope
name	string	query	False	Filter by name
licenses.serial_number	string	query	False	Filter by licenses.serial_number
licenses.start_time	string	query	False	Filter by licenses.start_time
licenses.expiry_time	string	query	False	Filter by licenses.expiry_time
licenses.capacity.used_size	integer	query	False	Filter by licenses.capacity.used_size
licenses.capacity.maximum_size	integer	query	False	Filter by licenses.capacity.maximum_size
licenses.compliance.state	string	query	False	Filter by licenses.compliance.state
licenses.evaluation	boolean	query	False	Filter by licenses.evaluation
licenses.active	boolean	query	False	Filter by licenses.active
licenses.owner	string	query	False	Filter by licenses.owner
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "licenses": [
        {
          "capacity": {
            "maximum_size": 0,
            "used_size": 0
          },
          "compliance": {
            "state": "compliant"
          },
          "expiry_time": "2019-03-02T19:00:00Z",
          "owner": "cluster1",
          "serial_number": "123456789",
          "start_time": "2019-02-02T19:00:00Z"
        }
      ],
      "name": "NFS",
      "scope": "string",
      "state": "compliant"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

capacity

Name	Type	Description
maximum_size	integer	Licensed capacity size (in bytes) that can be used.
used_size	integer	Capacity that is currently used (in bytes).

compliance

Name	Type	Description
state	string	Compliance state of the license.

licenses

Name	Type	Description
active	boolean	A flag indicating whether the license is currently being enforced.
capacity	capacity	
compliance	compliance	
evaluation	boolean	A flag indicating whether the license is in evaluation mode.

Name	Type	Description
expiry_time	string	Date and time when the license expires.
owner	string	Cluster, node or license manager that owns the license.
serial_number	string	Serial number of the license.
start_time	string	Date and time when the license starts.

records

Name	Type	Description
_links	_links	
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Install one or more feature licenses

POST /cluster/licensing/licenses

Introduced In: 9.6

Installs one or more feature licenses.

Required properties

- `keys` - Array containing a list of NLF or 26-character license keys.

Related ONTAP commands

- `system license add`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
keys	array[string]	
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

Example request

```
{
  "keys": [
    "AMEPOSOIKLKGEEEEEDGNDEKSJDE"
  ],
  "licenses": [
    {
      "capacity": {
        "maximum_size": 0,
        "used_size": 0
      },
      "compliance": {
        "state": "compliant"
      },
      "expiry_time": "2019-03-02T19:00:00Z",
      "owner": "cluster1",
      "serial_number": "123456789",
      "start_time": "2019-02-02T19:00:00Z"
    }
  ],
  "name": "NFS",
  "scope": "string",
  "state": "compliant"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "records": [
    {
      "keys": [
        "AMEPOSOIKLKGEEDGNDEKSJDE"
      ],
      "licenses": [
        {
          "capacity": {
            "maximum_size": 0,
            "used_size": 0
          },
          "compliance": {
            "state": "compliant"
          },
          "expiry_time": "2019-03-02T19:00:00Z",
          "owner": "cluster1",
          "serial_number": "123456789",
          "start_time": "2019-02-02T19:00:00Z"
        }
      ],
      "name": "NFS",
      "scope": "string",
      "state": "compliant"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1115117	Generic licensing error
1115122	No cluster serial number found
1115124	No node serial number found
1115130	No license code was provided
1115131	Installation of the license failed

Error Code	Description
1115132	License already exists on system
1115134	Serial number does not belong to node
1115141	License data is invalid
1115142	License signature is invalid
1115143	Internal error applying the requested license
1115152	License does not apply to the platform
1115154	Unable to retrieve cluster ID
1115155	Invalid cluster ID found
1115159	License is not in an acceptable format
1115160	License has already expired
1115164	Minimum ONTAP version requirements not met
1115179	FlexCache is not supported on this system
1115180	FlexCache is not supported on cloud systems
1115407	Capacity pool licenses cannot be installed directly
1115427	License is incompatible with capacity pools licensing mode
66846818	Failed to interpret FlexCache license information
66846821	FlexCache is not supported on cloud systems
66846822	Invalid FlexCache capacity information provided
655294464	Failed to extract license contents
655294465	License key is invalid
655294466	Serial number is invalid
655294467	Version number is invalid
655294468	Expired license
655294469	License does not apply to the platform
655294470	License does not apply to the product

Name	Type	Description
errors	array[error]	

Example error

```
{
  "errors": [
    {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    }
  ]
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

capacity

Name	Type	Description
maximum_size	integer	Licensed capacity size (in bytes) that can be used.
used_size	integer	Capacity that is currently used (in bytes).

compliance

Name	Type	Description
state	string	Compliance state of the license.

licenses

Name	Type	Description
active	boolean	A flag indicating whether the license is currently being enforced.
capacity	capacity	
compliance	compliance	
evaluation	boolean	A flag indicating whether the license is in evaluation mode.
expiry_time	string	Date and time when the license expires.
owner	string	Cluster, node or license manager that owns the license.
serial_number	string	Serial number of the license.

Name	Type	Description
start_time	string	Date and time when the license starts.

license_package

Name	Type	Description
keys	array[string]	
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

records

Name	Type	Description
keys	array[string]	
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a license

DELETE /cluster/licensing/licenses/{name}

Introduced In: 9.6

Deletes a license.

Related ONTAP commands

- `system license delete`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Name of the license package to delete.
serial_number	string	query	True	Serial number of the license to delete.

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
525028	Error during volume limit check, cannot remove license

Error Code	Description
525029	Current volume use will exceed limits if license is removed
1115137	Cluster license requires a base license to be installed
1115144	Cloud licenses cannot be deleted
1115178	A tier license that is still in use cannot be deleted
1115213	License is still in use and cannot be removed
1115406	Capacity pool licenses cannot be deleted
66846823	A FlexCache license that is still in use cannot be deleted

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a license package

GET /cluster/licensing/licenses/{name}

Introduced In: 9.6

Retrieves a specific license package.



By default, the GET method only returns licensed packages. You must provide the following query "state=unlicensed" to retrieve unlicensed packages.

Related ONTAP commands

- `system license show`
- `system license show-status`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Name of the license package.
state	string	query	False	Filter by state

Name	Type	In	Required	Description
scope	string	query	False	Filter by scope
name	string	query	False	Filter by name
licenses.serial_number	string	query	False	Filter by licenses.serial_number
licenses.start_time	string	query	False	Filter by licenses.start_time
licenses.expiry_time	string	query	False	Filter by licenses.expiry_time
licenses.capacity.used_size	integer	query	False	Filter by licenses.capacity.used_size
licenses.capacity.maximum_size	integer	query	False	Filter by licenses.capacity.maximum_size
licenses.compliance.state	string	query	False	Filter by licenses.compliance.state
licenses.evaluation	boolean	query	False	Filter by licenses.evaluation
licenses.active	boolean	query	False	Filter by licenses.active
licenses.owner	string	query	False	Filter by licenses.owner
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
licenses	array[licenses]	Installed licenses of the package.
name	string	Name of the license.
scope	string	Scope of the license.
state	string	Summary state of package based on all installed licenses.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "licenses": [
    {
      "capacity": {
        "maximum_size": 0,
        "used_size": 0
      },
      "compliance": {
        "state": "compliant"
      },
      "expiry_time": "2019-03-02T19:00:00Z",
      "owner": "cluster1",
      "serial_number": "123456789",
      "start_time": "2019-02-02T19:00:00Z"
    }
  ],
  "name": "NFS",
  "scope": "string",
  "state": "compliant"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

capacity

Name	Type	Description
maximum_size	integer	Licensed capacity size (in bytes) that can be used.
used_size	integer	Capacity that is currently used (in bytes).

compliance

Name	Type	Description
state	string	Compliance state of the license.

licenses

Name	Type	Description
active	boolean	A flag indicating whether the license is currently being enforced.
capacity	capacity	
compliance	compliance	
evaluation	boolean	A flag indicating whether the license is in evaluation mode.
expiry_time	string	Date and time when the license expires.
owner	string	Cluster, node or license manager that owns the license.

Name	Type	Description
serial_number	string	Serial number of the license.
start_time	string	Date and time when the license starts.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster mediators

Cluster mediators endpoint overview

Overview

You can use this API to connect to a mediator, delete or get the status of the mediator. The GET operation returns the status of the mediator. The DELETE operation deletes the mediator. The POST operation creates the mediator.

Examples

GET - MetroCluster status and configuration

```
[[ID6097d46eed9b3513ba79722e8a2cd4a4]]
= Retrieve ONTAP Mediators configured in the cluster
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/cluster/mediators`#
```

Introduced In: 9.8

Retrieves mediators configured in the cluster.

== Parameters

```
[cols=5*,options=header]  
|==
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|ip_address  
|string  
|query  
|False  
a|Filter by ip_address
```

```
|port  
|integer  
|query  
|False  
a|Filter by port
```

```
|peer_cluster.name  
|string  
|query  
|False  
a|Filter by peer_cluster.name
```

```
|peer_cluster.uuid  
|string  
|query  
|False  
a|Filter by peer_cluster.uuid
```

```
|reachable
|boolean
|query
|False
a|Filter by reachable
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```



```

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|

|records
|array[link:#records[records]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {

```

```

    "href": "/api/resourcelink"
  },
  "self": {
    "href": "/api/resourcelink"
  }
},
"records": [
  {
    "ip_address": "10.10.10.7",
    "peer_cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster2",
      "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
    },
    "port": "31784",
    "reachable": 1,
    "uuid": "string"
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]

```

```

=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

```

```

|self
|link:#href[href]
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|self
|link:#href[href]
a|

```

```

|===

```

```

[#peer_cluster]
[.api-collapsible-fifth-title]
peer_cluster

```

The peer cluster that the mediator service is used for.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|_links

```

```

|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#records]
[.api-collapsible-fifth-title]
records

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ip_address
|string
a|The IP address of the mediator.

|peer_cluster
|link:#peer_cluster[peer_cluster]
a|The peer cluster that the mediator service is used for.

|port
|integer
a|The REST server's port number on the mediator.

|reachable
|boolean
a|Indicates the connectivity status of the mediator.

|uuid
|string
a|The unique identifier for the mediator service.

```

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDde9035f5e75443e31de37d5122274a91]]
= Create and connect an ONTAP Mediator

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/cluster/mediators`#

*Introduced In:* 9.8

Creates and connect a mediator.

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the

```

default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

- * Default value: 1
- * Max value: 120
- * Min value: 0

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

- * Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|ip_address

|string

a|The IP address of the mediator.

|password

|string

a|The password used to connect to the REST server on the mediator.

|peer_cluster

|link:#peer_cluster[peer_cluster]

a|The peer cluster that the mediator service is used for.

|port


```

|integer
a|The REST server's port number on the mediator.

|reachable
|boolean
a|Indicates the connectivity status of the mediator.

|user
|string
a|The username used to connect to the REST server on the mediator.

|uuid
|string
a|The unique identifier for the mediator service.

|===

```

```

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "ip_address": "10.10.10.7",
  "password": "mypassword",
  "peer_cluster": {
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebef2"
  },
  "port": "31784",
  "reachable": 1,
  "user": "myusername",
  "uuid": "string"
}
====

```

== Response

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Error
```

Status: Default

ONTAP Error Response codes

```
|===
| Error code | Description

| 13369351
| Update to mediator failed. Reason: does not authorized for that command.
Check that the peer cluster and mediator are reachable.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
====  
[#href]  
[.api-collapsible-fifth-title]  
href
```

```
[cols=3*,options=header]  
|====  
|Name  
|Type  
|Description
```

```
|href  
|string  
a|
```

```
|====
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links  
[#peer_cluster]  
[.api-collapsible-fifth-title]  
peer_cluster
```

The peer cluster that the mediator service is used for.

```
[cols=3*,options=header]  
|====  
|Name  
|Type  
|Description
```

```
|name  
|string  
a|
```

```
|uuid  
|string  
a|
```

```
|====
```

```
[#mediator]  
[.api-collapsible-fifth-title]
```

mediator

Mediator information

[cols=3*,options=header]

|===

|Name

|Type

|Description

|ip_address

|string

a|The IP address of the mediator.

|password

|string

a|The password used to connect to the REST server on the mediator.

|peer_cluster

|link:#peer_cluster[peer_cluster]

a|The peer cluster that the mediator service is used for.

|port

|integer

a|The REST server's port number on the mediator.

|reachable

|boolean

a|Indicates the connectivity status of the mediator.

|user

|string

a|The username used to connect to the REST server on the mediator.

|uuid

|string

a|The unique identifier for the mediator service.

|===

```

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID87c3c786ef37f1aadbd0c77c30ef3d01]]
= Delete an ONTAP Mediator

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/cluster/mediators/{uuid}`#

*Introduced In:* 9.8

Deletes the mediator.

== Parameters

```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|* format: uuid
```

```
|return_timeout
```

```
|integer
```

```
|query
```

```
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|ip_address
```

```
|string
```

a|The IP address of the mediator.

|password

|string

a|The password used to connect to the REST server on the mediator.

|peer_cluster

|link:#peer_cluster[peer_cluster]

a|The peer cluster that the mediator service is used for.

|port

|integer

a|The REST server's port number on the mediator.

|reachable

|boolean

a|Indicates the connectivity status of the mediator.

|user

|string

a|The username used to connect to the REST server on the mediator.

|uuid

|string

a|The unique identifier for the mediator service.

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "ip_address": "10.10.10.7",
  "password": "mypassword",
  "peer_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
```



```
    }  
  },  
  "name": "cluster2",  
  "uuid": "ebe27c49-1adf-4496-8335-ab862aebbf2"  
},  
"port": "31784",  
"reachable": 1,  
"user": "myusername",  
"uuid": "string"  
}  
====  
  
== Response
```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
=====

== Error

```

Status: Default

ONTAP Error Response codes

```

|===
| Error code | Description

| 13369377
| Mediator field "mediator.id" does not exist.
|===

```

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type

```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#peer_cluster]
```

```
[.api-collapsible-fifth-title]
```

```
peer_cluster
```

The peer cluster that the mediator service is used for.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```

|uuid
|string
a|

|===

[#mediator]
[.api-collapsible-fifth-title]
mediator

Mediator information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ip_address
|string
a|The IP address of the mediator.

|password
|string
a|The password used to connect to the REST server on the mediator.

|peer_cluster
|link:#peer_cluster[peer_cluster]
a|The peer cluster that the mediator service is used for.

|port
|integer
a|The REST server's port number on the mediator.

|reachable
|boolean
a|Indicates the connectivity status of the mediator.

|user
|string

```

```

a|The username used to connect to the REST server on the mediator.

|uuid
|string
a|The unique identifier for the mediator service.

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```

[[ID554cab63dce2dc98d92f7bf2614be51d]]
= Retrieve ONTAP Mediator state and configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/cluster/mediators/{uuid}`#

*Introduced In:* 9.8

Retrieves the mediator state and configuration.

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|* format: uuid

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type

```



```

|Description

|ip_address
|string
a|The IP address of the mediator.

|peer_cluster
|link:#peer_cluster[peer_cluster]
a|The peer cluster that the mediator service is used for.

|port
|integer
a|The REST server's port number on the mediator.

|reachable
|boolean
a|Indicates the connectivity status of the mediator.

|uuid
|string
a|The unique identifier for the mediator service.

```

```

|===

```

```

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "ip_address": "10.10.10.7",
  "peer_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebbf2"
  },
  "port": "31784",
  "reachable": 1,

```

```

    "uuid": "string"
  }
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#peer_cluster]
[.api-collapsible-fifth-title]
peer_cluster

The peer cluster that the mediator service is used for.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====


[[ID5135181ec324e97eec9691ee2ba32e85]]
= Retrieve historical performance metrics for the cluster


[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/cluster/metrics`#


*Introduced In:* 9.6


Retrieves historical performance metrics for the cluster.


== Parameters


[cols=5*,options=header]
|===


|Name

```

```
|Type
|In
|Required
|Description

|latency.total
|integer
|query
|False
a|Filter by latency.total
```

```
|latency.read
|integer
|query
|False
a|Filter by latency.read
```

```
|latency.other
|integer
|query
|False
a|Filter by latency.other
```

```
|latency.write
|integer
|query
|False
a|Filter by latency.write
```

```
|iops.total
|integer
|query
|False
a|Filter by iops.total
```

```
|iops.read
|integer
|query
|False
a|Filter by iops.read
```

```
|iops.other
|integer
|query
|False
a|Filter by iops.other
```

```
|iops.write
|integer
|query
|False
a|Filter by iops.write
```

```
|status
|string
|query
|False
a|Filter by status
```

```
|duration
|string
|query
|False
a|Filter by duration
```

```
|timestamp
|string
|query
|False
a|Filter by timestamp
```

```
|throughput.total
|integer
|query
|False
a|Filter by throughput.total
```

```
|throughput.read
|integer
|query
|False
a|Filter by throughput.read
```

```

|throughput.other
|integer
|query
|False
a|Filter by throughput.other

|throughput.write
|integer
|query
|False
a|Filter by throughput.write

|interval
|string
|query
|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, or 1y.
The period for each time range is specified as follows:

* 1h: Metrics over the most recent hour sampled over 15 seconds.
* 1d: Metrics over the most recent day sampled over 5 minutes.
* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|fields

```



```

|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|num_records
|integer
a|Number of records
```

```
|records
|array[link:#records[records]]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
```

```

        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

|===

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

|===

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#records]
[.api-collapsible-fifth-title]
records

Performance numbers, such as IOPS latency and throughput.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#error_arguments]

```



```
[.api-collapsible-fifth-title]
```

```
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
:leveloffset: -1
```

```
= View and manage MetroCluster configurations
```

```
:leveloffset: +1
```

```
[[ID437eb998b66e15f517531651cd11bc7c]]
= Cluster MetroCluster endpoint overview
```

```
= Overview
```

You can use this API to create, perform operations, and retrieve relevant information pertaining to MetroCluster. The GET operation fetches MetroCluster status and configuration parameters for the local and partner cluster. The PATCH operation executes a switchover or switchback operation. The POST request can be used to setup a MetroCluster.

```
== Creating a MetroCluster
```

A new MetroCluster can be set up by issuing a POST to /cluster/metrocluster. Parameters are provided in the body of the POST request.

```
=== Fields used for setting up a MetroCluster configuration
```

The fields used for MetroCluster APIs are either required or optional and are described as follows:

```
=== Required configuration fields
```

These fields are always required for any POST /cluster/metrocluster request.

* `partner_cluster.name` - Specifies the partner cluster name to which cluster peering has been established.

* `dr_pairs` - Specifies local and DR partner node pairs. Each pair uniquely identifies a DR group.

=== Optional configuration fields

This field is used to set up additional components in a MetroCluster configuration.

* `mediator.+++` - Specifies mediator parameters. If Mediator Assisted Unplanned Switchover (MAUSO) functionality is required, then a mediator should be configured.

=== Polling the setup job

After a successful POST /cluster/metrocluster is issued, an HTTP status code of 202 (Accepted) is returned along with a job UUID and a link in the body of the response. The setup job continues asynchronously and can be monitored by using the job UUID and the /cluster/jobs API. The "message" field in the response of the GET /cluster/jobs/{uuid} request shows the current step in the job, and the "state" field shows the overall state of the job.

'''

== Examples

== Setting up a 4-node MetroCluster

This example shows the POST body when setting up a 4-node MetroCluster along with a mediator. It is required that cluster peering be established between two clusters, in this example, site "mcc_siteA" and "mcc_siteB" before issuing the POST request. Nodes "node-a" and "node-b" are HA partners and part of the local cluster "mcc_siteA", whereas nodes "node-c" and "node-d" are HA partners in the partner cluster "mcc_siteB". Specifying a single DR pairing of "node-a" and "node-c" is sufficient to identify a DR group -- "node-a" and "node-c" will be designated primary DR partners ("node-b" and "node-d" too). "node-d" will then be designated auxiliary partner of "node-a". Once the MetroCluster configuration has been completed, and since mediator parameters have been provided, the mediator will be setup and MAUSO enabled.

```

# API
/api/cluster/metrocluster
----

=== POST body included from file

----
mcc_post_body.txt:
{
  "partner_cluster" : {
    "name": "mcc_siteB"
  },
  "dr_pairs" : [
    {
      "node" : {
        "name" : "node-a"
      },
      "partner" : {
        "name" : "node-c"
      }
    }
  ],
  "mediator" : {
    "ip_address" : "1.2.3.4",
    "user" : "mcc_mediator",
    "password" : "openMediator"
  }
}

curl -X POST https://<mgmt-ip>/api/cluster/metrocluster -d
"@mcc_post_body.txt"
----

=== Inline POST body

----
curl -X POST https://<mgmt-ip>/api/cluster/metrocluster -d
'{"partner_cluster" : {"name": "mcc_siteB" }, "dr_pairs" : [{"node" :
{"name" : "node-a" }, "partner" : {"name" : "node-c" } ]}, "mediator" :
{"ip_address" : "1.2.3.4", "user" : "mcc_mediator" , "password" :
"openMediator" } }'
----

=== POST Response

----
HTTP/1.1 202 Accepted

```

```
Date: Thu, 09 Jan 2020 20:38:05 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/cluster/metrocluster
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "f23abdb-331f-11ea-acd3-005056a708b2",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f23abdb-331f-11ea-acd3-005056a708b2"
      }
    }
  }
}
-----
```

=== Monitoring the job progress

Use the link provided in the response to the POST request to fetch information for the setup job.

==== Request

```
-----
curl -X GET https://<mgmt-ip>/api/cluster/jobs/f23abdb-331f-11ea-acd3-005056a708b2
-----
```

==== Job status response

The following is an example of the job status response returned by the running MetroCluster setup job:

```
-----
HTTP/1.1 200 OK
Date: Thu, 09 Jan 2020 20:40:20 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 373
Content-Type: application/hal+json
{
  "uuid": "f23abdb-331f-11ea-acd3-005056a708b2",
```

```

"description": "POST /api/cluster/metrocluster",
"state": "running",
"message": "Checking remote storage pool",
"code": 2432844,
"start_time": "2020-01-09T15:38:08-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f23abbdb-331f-11ea-acd3-005056a708b2"
  }
}

```

=== Completion message

This is the final update message from the setup job indicating completion.

```

----
{
  "uuid": "f23abbdb-331f-11ea-acd3-005056a708b2",
  "description": "POST /api/cluster/metrocluster",
  "state": "running",
  "message": "MetroCluster setup is complete",
  "code": 2432849,
  "start_time": "2020-01-09T15:38:08-05:00",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f23abbdb-331f-11ea-acd3-005056a708b2"
    }
  }
}

```

=== Final status of a successful MetroCluster setup workflow

When the setup job completes, the 'end_time' field is populated, and the 'state' and 'message' fields report the final status.

```

----
HTTP/1.1 200 OK
Date: Thu, 09 Jan 2020 20:43:54 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 360
Content-Type: application/hal+json
{

```

```
"uuid": "f23abdbd-331f-11ea-acd3-005056a708b2",
"description": "POST /api/cluster/metrocluster",
"state": "success",
"message": "success",
"code": 0,
"start_time": "2020-01-09T15:38:08-05:00",
"end_time": "2020-01-09T15:43:50-05:00",
"_links": {
  "self": {
    "href": "/api/cluster/jobs/f23abdbd-331f-11ea-acd3-005056a708b2"
  }
}
}
```

==== Retrieving the MetroCluster configuration after completion of the POST request

==== Request

```
-----
curl -X GET https://<mgmt-ip>/api/cluster/metrocluster
-----
```

==== Response

```
-----
HTTP/1.1 200 OK
Date: Thu, 09 Jan 2020 20:49:40 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 849
Content-Type: application/hal+json
{
  "local": {
    "configuration_state": "configured",
    "periodic_check_enabled": true,
    "mode": "normal",
    "partner_cluster_reachable": true,
    "cluster": {
      "name": "mcc_siteA",
      "uuid": "4294c4f2-30e2-11ea-8cac-005056a708b2",
      "_links": {
        "self": {
          "href": "/api/cluster/4294c4f2-30e2-11ea-8cac-005056a708b2"
        }
      }
    }
  }
}
```

```

    }
  }
},
"remote": {
  "configuration_state": "configured",
  "periodic_check_enabled": true,
  "mode": "normal",
  "cluster": {
    "name": "mcc_siteB",
    "uuid": "4207c6a5-30e2-11ea-be25-005056a7dc84",
    "_links": {
      "self": {
        "href": "/api/cluster/4207c6a5-30e2-11ea-be25-005056a7dc84"
      }
    }
  }
},
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster"
  }
}
}
}
-----

```

=== Retrieving information about the nodes in a MetroCluster configuration

==== Request

```

-----
curl -X GET https://<mgmt-ip>/api/cluster/metrocluster/nodes
-----

```

==== Response

```

-----
HTTP/1.1 200 OK
Date: Fri, 10 Jan 2020 02:26:20 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/hal+json
Transfer-Encoding: chunked
{
"records": [

```



```

{
  "dr_group_id": 1,
  "cluster": {
    "name": "mcc_siteA",
    "uuid": "4294c4f2-30e2-11ea-8cac-005056a708b2",
    "_links": {
      "self": {
        "href": "/api/cluster/4294c4f2-30e2-11ea-8cac-005056a708b2"
      }
    }
  },
  "node": {
    "name": "node-a",
    "uuid": "1e6b0137-30dd-11ea-82ba-005056a7c78a",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/1e6b0137-30dd-11ea-82ba-005056a7c78a"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/cluster/metrocluster/nodes/1e6b0137-30dd-11ea-82ba-005056a7c78a"
    }
  }
},
{
  "dr_group_id": 1,
  "cluster": {
    "name": "mcc_siteA",
    "uuid": "4294c4f2-30e2-11ea-8cac-005056a708b2",
    "_links": {
      "self": {
        "href": "/api/cluster/4294c4f2-30e2-11ea-8cac-005056a708b2"
      }
    }
  },
  "node": {
    "name": "node-b",
    "uuid": "1e57ba22-30dd-11ea-8b19-005056a708b2",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/1e57ba22-30dd-11ea-8b19-005056a708b2"
      }
    }
  }
}

```

```

    }
  },
  "_links": {
    "self": {
      "href": "/api/cluster/metrocluster/nodes/1e57ba22-30dd-11ea-8b19-005056a708b2"
    }
  }
},
{
  "dr_group_id": 1,
  "cluster": {
    "name": "mcc_siteB",
    "uuid": "4207c6a5-30e2-11ea-be25-005056a7dc84",
    "_links": {
      "self": {
        "href": "/api/cluster/4207c6a5-30e2-11ea-be25-005056a7dc84"
      }
    }
  },
  "node": {
    "name": "node-c",
    "uuid": "1e563efc-30dd-11ea-a9d3-005056a71573",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/1e563efc-30dd-11ea-a9d3-005056a71573"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/cluster/metrocluster/nodes/1e563efc-30dd-11ea-a9d3-005056a71573"
    }
  }
},
{
  "dr_group_id": 1,
  "cluster": {
    "name": "mcc_siteB",
    "uuid": "4207c6a5-30e2-11ea-be25-005056a7dc84",
    "_links": {
      "self": {
        "href": "/api/cluster/4207c6a5-30e2-11ea-be25-005056a7dc84"
      }
    }
  }
}

```

```

    }
  }
},
"node": {
  "name": "node-d",
  "uuid": "1e400aa4-30dd-11ea-adeb-005056a7dc84",
  "_links": {
    "self": {
      "href": "/api/cluster/nodes/1e400aa4-30dd-11ea-adeb-005056a7dc84"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/nodes/1e400aa4-30dd-11ea-adeb-005056a7dc84"
  }
}
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/nodes"
  }
}
}
}
----

'''

```

=== Retrieving MetroCluster status and configuration information

```

----
GET https://<mgmt-ip>/api/cluster/metrocluster
{
  "local": {
    "configuration_state": "configured",
    "periodic_check_enabled": true,
    "mode": "normal",
    "cluster": {
      "name": "cluster1",
      "uuid": "bbc00ca3-8d81-11e9-b5a9-005056826931",
      "_links": {
        "self": {

```

```

        "href": "/api/cluster/bbc00ca3-8d81-11e9-b5a9-
005056826931"
    }
}
},
"remote": {
    "configuration_state": "configured",
    "periodic_check_enabled": true,
    "mode": "normal",
    "cluster": {
        "name": "cluster3",
        "uuid": "ce2cf803-8d81-11e9-87db-00505682cecf",
        "_links": {
            "self": {
                "href": "/api/cluster/ce2cf803-8d81-11e9-87db-
00505682cecf"
            }
        }
    }
},
"_links": {
    "self": {
        "href": "/api/cluster/metrocluster"
    }
}
}
----
'''

```

=== Initiating a switchover or switchback command using PATCH

PATCH is used to initiate a variety of operations by specifying one of the following values in the "action" parameter:

- * `switchover` - Initiates an Unplanned Switchover (USO).
- * `negotiated_switchover` - Indicates that an Negotiated switchover (NSO) is to be performed.
- * `negotiated_switchover_simulate` - Provides validation in preparation for NSO but does not perform the operation.
- * `switchback` - Indicates that a switchback is to be performed.
- * `switchback_simulate` - Provides validation for switchback but does not commit the operation.

==== PATCH Switchover example

```

-----
PATCH https://<mgmt-ip>/api/cluster/metrocluster?action=switchover
{
  "job": {
    "uuid": "70e54274-57ee-11e9-aa33-005056820b99",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/70e54274-57ee-11e9-aa33-
005056820b99"
      }
    }
  }
}
}
-----

```

This returns a job UUID. A subsequent GET for this job should return the following:

```

-----
GET https://<mgmt-ip>/api/cluster/jobs/70e54274-57ee-11e9-aa33-
005056820b99
{
  "uuid": "70e54274-57ee-11e9-aa33-005056820b99",
  "description": "MetroCluster Switchover Job",
  "state": "success",
  "message": "Complete: Switchover is successful.",
  "code": 0,
  "start_time": "2019-04-05T15:02:02-07:00",
  "end_time": "2019-04-05T15:02:30-07:00",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/70e54274-57ee-11e9-aa33-005056820b99"
    }
  }
}
}
-----

```

==== PATCH Switchback example:

```

-----
PATCH https://<mgmt-ip>/api/cluster/metrocluster?action=switchback
{
  "job": {
    "uuid": "a62714cc-57ec-11e9-aa33-005056820b99",
    "_links": {
      "self": {

```

```

        "href": "/api/cluster/jobs/a62714cc-57ec-11e9-aa33-005056820b99"
      }
    }
  }
}
-----

```

This returns a job UUID with a link to the job. A subsequent GET for this job UUID can be used to retrieve the completion status of the operation:

```

-----
GET https://<mgmt-ip>/api/cluster/jobs/a62714cc-57ec-11e9-aa33-005056820b99
{
  "uuid": "a62714cc-57ec-11e9-aa33-005056820b99",
  "description": "MetroCluster Switchback Job",
  "state": "success",
  "message": "Complete: Switchback is successful.",
  "code": 0,
  "start_time": "2019-04-05T14:49:12-07:00",
  "end_time": "2019-04-05T14:50:12-07:00",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/a62714cc-57ec-11e9-aa33-005056820b99"
    }
  }
}
-----

```

```

[[IDfce2c449271f8f605e828113624419d5]]
= Retrieve MetroCluster status and configuration details

```

```

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/cluster/metrocluster`#

```

Introduced In: 9.8

Retrieves MetroCluster status and configuration details.

== Related ONTAP commands * `metrocluster show` * `metrocluster node show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|max_records
|integer
|query
|False

a|Limit the number of records returned.

|return_records
|boolean
|query
|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction.

Default direction is 'asc' for ascending.

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|enabled
|boolean
a|

|local
|link:#local[local]
a|

|remote
|link:#remote[remote]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
```



```

    "self": {
      "href": "/api/resourcelink"
    },
    "local": {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "configuration_state": "string",
      "mode": "string"
    },
    "remote": {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "configuration_state": "string",
      "mode": "string"
    }
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

2425734	
---------	--

An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support.

```

|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

Local node of the DR Group.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the node.

```

```
|===
```

```
[#partner]  
[.api-collapsible-fifth-title]  
partner
```

Partner node of the DR Group.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|Name of the node.
```

```
|===
```

```
[#dr_pairs]  
[.api-collapsible-fifth-title]  
dr_pairs
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|node  
|link:#node[node]  
a|Local node of the DR Group.
```

```
|partner  
|link:#partner[partner]  
a|Partner node of the DR Group.
```

```
|===
```

```

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#local]
[.api-collapsible-fifth-title]
local

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|configuration_state
|string
a|Indicates the state of the local cluster configuration.

|mode
|string
a|Specifies the mode of operation of the local cluster.

|partner_cluster_reachable
|boolean
a|Specifies whether the partner cluster is reachable from the local
cluster.

|periodic_check_enabled
|boolean
a|Indicates whether or not a periodic check is enabled on the local
cluster.

|===

[#peer_cluster]
[.api-collapsible-fifth-title]
peer_cluster

The peer cluster that the mediator service is used for.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#mediator]
[.api-collapsible-fifth-title]
mediator
```

Mediator information

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|ip_address
|string
a|The IP address of the mediator.
```

```
|peer_cluster
|link:#peer_cluster[peer_cluster]
a|The peer cluster that the mediator service is used for.
```

```
|port
|integer
a|The REST server's port number on the mediator.
```

```
|reachable
|boolean
a|Indicates the connectivity status of the mediator.
```

```
|uuid
|string
a|The unique identifier for the mediator service.
```

```
|===
```

```
[#partner_cluster]
[.api-collapsible-fifth-title]
partner_cluster
```

Partner cluster information.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|Name of the partner cluster.
```

```
|===
```

```
[#remote]
```

```
[.api-collapsible-fifth-title]
```

```
remote
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cluster
```

```
|link:#cluster[cluster]
```

```
a|
```

```
|configuration_state
```

```
|string
```


a|Indicates the state of the remote cluster configuration.

|mode

|string

a|Specifies the mode of operation of the remote cluster.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID5bdd8db7d63994d04ab996dbf5571115]]
= Initiate a switchover, heal, or switchback operation

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/cluster/metrocluster`#

*Introduced In:* 9.8

Initiates a switchover or switchback operation.

== Related ONTAP commands * `metrocluster switchover` * `metrocluster
switchback`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In

```

```

|Required
|Description

|action
|string
|query
|False
a|Action to perform on the MetroCluster.

* enum: ["switchover", "negotiated_switchover",
"negotiated_switchover_simulate", "switchback", "switchback_simulate"]

|return_timeout
|integer
|query
|False
a|Timeout in seconds for the call.

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds. This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job. If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 1
* Max value: 120
* Min value: 0

|===

```

== Response

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
=====

== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description

| 2425734
| An internal error occurred. Wait a few minutes, and try the operation
again. For further assistance, contact technical support.
|===

== Definitions
```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[ID31c59702be4c22e3f5e1b25e4d89fa21]]

= Set up a MetroCluster configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/cluster/metrocluster`#

Introduced In: 9.8

Sets up a MetroCluster.

== Required properties

* `partner_cluster.name`

* `dr_pairs`

== Recommended optional properties

* `mediator.+++`

== Learn more

* xref:{relative_path}cluster_metrocluster_endpoint_overview.html[DOC /cluster/metrocluster]

== Related ONTAP commands

* `metrocluster configuration-settings dr-group create`

* `metrocluster configuration-settings interface create`

* `metrocluster configuration-settings connection connect`

* `metrocluster configuration-settings mediator add`

* `storage aggregate create`

* `storage aggregate mirror`

* `metrocluster configure`

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|return_timeout
```

```
|integer
```

```
|query
```

```
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

```
|return_records
```

```
|boolean
```

```
|query
```

```
|False
```

a|The default is false. If set to true, the records are returned.

* Default value:

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```
|dr_pairs
|array[link:#dr_pairs[dr_pairs]]
a|DR Pairs to create as part of a MetroCluster configure.
```

```
|enabled
|boolean
a|
```

```
|local
|link:#local[local]
a|
```

```
|mediator
|link:#mediator[mediator]
a|Mediator information
```

```
|partner_cluster
|link:#partner_cluster[partner_cluster]
a|Partner cluster information.
```

```
|remote
|link:#remote[remote]
a|
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "dr_pairs": [
    {
      "node": {
        "name": "nodeA"
      },
      "partner": {
        "name": "nodeB"
      }
    }
  ],
  "local": {
    "cluster": {
```

```

        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "configuration_state": "string",
    "mode": "string"
},
"mediator": {
    "ip_address": "10.10.10.7",
    "password": "mypassword",
    "peer_cluster": {
        "name": "cluster2",
        "uuid": "ebe27c49-1adf-4496-8335-ab862aebef2"
    },
    "port": "31784",
    "reachable": 1,
    "user": "myusername",
    "uuid": "string"
},
"partner_cluster": {
    "name": "cluster2"
},
"remote": {
    "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "configuration_state": "string",
    "mode": "string"
}
}
====

== Response

```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
=====

== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description

| 2425734
| An internal error occurred. Wait a few minutes, and try the operation
again. For further assistance, contact technical support.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
====  
[#href]  
[.api-collapsible-fifth-title]  
href
```

```
[cols=3*,options=header]  
|====  
|Name  
|Type  
|Description
```

```
|href  
|string  
a|
```

```
|====
```

```
[#self_link]  
[.api-collapsible-fifth-title]  
self_link  
[#node]  
[.api-collapsible-fifth-title]  
node
```

Local node of the DR Group.

```
[cols=3*,options=header]  
|====  
|Name  
|Type  
|Description
```

```
|name  
|string  
a|Name of the node.
```

```
|====
```

```
[#partner]  
[.api-collapsible-fifth-title]  
partner
```

Partner node of the DR Group.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the node.

|===

[#dr_pairs]
[.api-collapsible-fifth-title]
dr_pairs

[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|Local node of the DR Group.

|partner
|link:#partner[partner]
a|Partner node of the DR Group.

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#local]
[.api-collapsible-fifth-title]
local

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|configuration_state
|string
a|Indicates the state of the local cluster configuration.

|mode
|string
a|Specifies the mode of operation of the local cluster.

|partner_cluster_reachable
|boolean
a|Specifies whether the partner cluster is reachable from the local
cluster.

|periodic_check_enabled

```

```
|boolean
a|Indicates whether or not a periodic check is enabled on the local
cluster.
```

```
|===
```

```
[#peer_cluster]
[.api-collapsible-fifth-title]
peer_cluster
```

The peer cluster that the mediator service is used for.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#mediator]
[.api-collapsible-fifth-title]
mediator
```

Mediator information

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|ip_address
|string
```

a|The IP address of the mediator.

|password

|string

a|The password used to connect to the REST server on the mediator.

|peer_cluster

|link:#peer_cluster[peer_cluster]

a|The peer cluster that the mediator service is used for.

|port

|integer

a|The REST server's port number on the mediator.

|reachable

|boolean

a|Indicates the connectivity status of the mediator.

|user

|string

a|The username used to connect to the REST server on the mediator.

|uuid

|string

a|The unique identifier for the mediator service.

|===

[#partner_cluster]

[.api-collapsible-fifth-title]

partner_cluster

Partner cluster information.

[cols=3*,options=header]

|===

|Name

|Type


```

|Description

|name
|string
a|Name of the partner cluster.

|===

[#remote]
[.api-collapsible-fifth-title]
remote

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|configuration_state
|string
a|Indicates the state of the remote cluster configuration.

|mode
|string
a|Specifies the mode of operation of the remote cluster.

|===

[#metrocluster]
[.api-collapsible-fifth-title]
metrocluster

Holds MetroCluster status and configuration parameters for the local and
remote clusters. REST: /api/cluster/metrocluster

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|dr_pairs
|array[link:#dr_pairs[dr_pairs]]
a|DR Pairs to create as part of a MetroCluster configure.

|enabled
|boolean
a|

|local
|link:#local[local]
a|

|mediator
|link:#mediator[mediator]
a|Mediator information

|partner_cluster
|link:#partner_cluster[partner_cluster]
a|Partner cluster information.

|remote
|link:#remote[remote]
a|

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or

```

DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
= Display MetroCluster diagnostics
```

```
:leveloffset: +1
```

```
[[ID75c34bb14e6ba7078651fad49e8048dd]]
= Cluster MetroCluster diagnostics endpoint overview
```

```
== Overview
```

You can use this API to initiate diagnostics and fetch the results of a completed diagnostic operation on a MetroCluster configuration.

```
=== Related ONTAP Commands
```

```
* `metrocluster check run`
```

```
[[IDc267dd3dab4aff40012968273e821ba0]]
```

= Retrieve diagnostic operation results for a MetroCluster configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/cluster/metrocluster/diagnostics`#

Introduced In: 9.8

Retrieves the results of a completed diagnostic operation for the MetroCluster configuration.

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|component
|string
|query
|False
a|MetroCluster diagnostic component name

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|components
|array[link:#metrocluster_diag_component[metrocluster_diag_component]]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "components": [
    {
      "aggregates": [
        {
          "aggregate": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "check": "string",
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "details": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "state": "string"
  ]
},
"clusters": [
  {
    "check": "string",
    "cluster": {
      "_links": {
```

```

        "self": {
            "href": "/api/resourcelink"
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "details": "string",
    "state": "string",
    "timestamp": "2016-03-10T14:35:16-08:00"
}
],
"connections": [
    {
        "cluster": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "config_error_info": "string",
        "config_state": "string",
        "destination_address": "10.10.10.7",
        "dr_group": {
            "id": 0
        },
        "home_port": "string",
        "mtu_error_details": "string",
        "node": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "partner_cluster": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },

```



```

        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "partner_node": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "partner_type": "string",
    "ping_error_details": "string",
    "relationship_type": "string",
    "source_address": "10.10.10.7",
    "state": "string",
    "storage_error_details": "string"
}
],
"details": "string",
"lifs": [
    {
        "check": "string",
        "cluster": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "details": "string",
        "lif": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "ip": {
                "address": "10.10.10.7"
            },
            "name": "lif1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
    },

```

```

    "state": "string",
    "vserver": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svml",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "nodes": [
    {
      "check": "string",
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "details": "string",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "state": "string"
  ]
},
"state": "string",
"timestamp": "2016-03-10T14:35:16-08:00",
"type": "string",
"volumes": [
  {
    "check": "string",
    "details": "string",
    "state": "string",
    "volume": {

```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "vserver": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
]
}
]
}
====
== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 2427132
| MetroCluster is not configured on this cluster.
|====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]

```

```

[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#aggregate]
[.api-collapsible-fifth-title]
aggregate

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links

```

```

|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#metrocluster_diag_component_aggr]
[.api-collapsible-fifth-title]
metrocluster_diag_component_aggr

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregate
|link:#aggregate[aggregate]
a|Aggregate

|check
|string
a|Type of check run on the component.

|cluster
|link:#cluster[cluster]
a|

|details
|string
a|Additional information or recovery steps to take.

|node

```

```

|link:#node[node]
a|

|state
|string
a|Status of the diagnostic operation on this component.

|===

[#metrocluster_diag_component_cluster]
[.api-collapsible-fifth-title]
metrocluster_diag_component_cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|check
|string
a|Type of check run on the component.

|cluster
|link:#cluster[cluster]
a|

|details
|string
a|Additional information or recovery steps to take.

|state
|string
a|Status of the diagnostic operation on this component.

|timestamp
|string
a|Time of the most recent check.

|===

```

```
[#dr_group]
[.api-collapsible-fifth-title]
dr_group
```

DR group reference.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|id
|integer
a|DR Group ID
```

```
|===
```

```
[#partner_cluster]
[.api-collapsible-fifth-title]
partner_cluster
```

Partner Cluster information

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
 |_links
|link:#_links[_links]
a|
```

```
 |name
|string
a|
```

```
 |uuid
|string
a|
```



```

|===

[#partner_node]
[.api-collapsible-fifth-title]
partner_node

Partner Node information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#metrocluster_diag_component_connection]
[.api-collapsible-fifth-title]
metrocluster_diag_component_connection

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|config_error_info
|string

```

a|Configuration error information.

|config_state

|string

a|Configuration state of the connection.

|destination_address

|string

a|Destination IP Address

|dr_group

|link:#dr_group[dr_group]

a|DR group reference.

|home_port

|string

a|Home port

|mtu_error_details

|string

a|Additional information or recovery steps associated with mtu size errors.

|node

|link:#node[node]

a|

|partner_cluster

|link:#partner_cluster[partner_cluster]

a|Partner Cluster information

|partner_node

|link:#partner_node[partner_node]

a|Partner Node information

|partner_type

|string

a|Partner's assigned role in a DR group.

```
|ping_error_details
|string
a|Additional information or recovery steps associated with ping errors.
```

```
|relationship_type
|string
a|This node's assigned role in a DR group.
```

```
|source_address
|string
a|Source IP Address
```

```
|state
|string
a|Status of diagnostic operation on this component.
```

```
|storage_error_details
|string
a|Additional information or recovery steps associated with storage errors.
```

```
|===
```

```
[#ip]
[.api-collapsible-fifth-title]
ip
```

IP information

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|address
|string
a|IPv4 or IPv6 address
```

```

|===

[#lif]
[.api-collapsible-fifth-title]
lif

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ip
|link:#ip[ip]
a|IP information


|name
|string
a|The name of the interface.


|uuid
|string
a|The UUID that uniquely identifies the interface.


|===

[#vserver]
[.api-collapsible-fifth-title]
vserver

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]

```

```

a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#metrocluster_diag_component_lif]
[.api-collapsible-fifth-title]
metrocluster_diag_component_lif

[cols=3*,options=header]
|===
|Name
|Type
|Description

|check
|string
a|Type of check run on the component.

|cluster
|link:#cluster[cluster]
a|

|details
|string
a|Additional information or recovery steps to take.

|lif
|link:#lif[lif]
a|

|state
|string
a|Status of the diagnostic operation on this component.

```

```

|vserver
|link:#vserver[vserver]
a|

|===

[#metrocluster_diag_component_node]
[.api-collapsible-fifth-title]
metrocluster_diag_component_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|check
|string
a|Type of check run on the component.

|cluster
|link:#cluster[cluster]
a|

|details
|string
a|Additional information or recovery steps to take.

|node
|link:#node[node]
a|

|state
|string
a|Status of the diagnostic operation on this component.

|===

[#volume]
[.api-collapsible-fifth-title]
volume

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the volume.

|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6

|===

[#metrocluster_diag_component_volume]
[.api-collapsible-fifth-title]
metrocluster_diag_component_volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|check
|string
a|Type of check run on the component.

|details
|string
a|Additional information or recovery steps to take.

```

```

|state
|string
a|Status of diagnostic operation on this component.

|volume
|link:#volume[volume]
a|

|vserver
|link:#vserver[vserver]
a|

|===

[#metrocluster_diag_component]
[.api-collapsible-fifth-title]
metrocluster_diag_component

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|array[link:#metrocluster_diag_component_aggr[metrocluster_diag_component_
aggr]]
a|This object contains details of the aggregate checks when the state
associated with the check is "warning".

|clusters
|array[link:#metrocluster_diag_component_cluster[metrocluster_diag_compone
nt_cluster]]
a|This object contains details of the cluster checks when the state
associated with the check is "warning".

|connections
|array[link:#metrocluster_diag_component_connection[metrocluster_diag_comp
onent_connection]]
a|This object contains details of the connections checks the state
associated with the check is "warning".

```



```

|details
|string
a|Additional information or recovery steps to take.

|lifs
|array[link:#metrocluster_diag_component_lif[metrocluster_diag_component_1
if]]
a|This object contains details of the interface checks when the state
associated with the check is "warning".

|nodes
|array[link:#metrocluster_diag_component_node[metrocluster_diag_component_
node]]
a|This object contains details of the node checks when the state
associated with the check is "warning".

|state
|string
a|Status of diagnostic operation for this component.

|timestamp
|string
a|Time of the most recent diagnostic operation for this component

|type
|string
a|MetroCluster diagnostic component name

|volumes
|array[link:#metrocluster_diag_component_volume[metrocluster_diag_componen
t_volume]]
a|This object contains details of the volume checks when the state
associated with the check is "warning".

|===

//end collapsible .Definitions block

```

====

[[IDda050c8ef318160eda8b2dbfd812aa5e]]

= Start MetroCluster diagnostics or set up a periodic diagnostic schedule

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/cluster/metrocluster/diagnostics`#

Introduced In: 9.8

Start a MetroCluster diagnostic operation or set up a schedule for the diagnostics to be run periodically.

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|schedule

|integer

|query

|False

a|Shows the minutes of every hour when a job runs. Setting this parameter schedules the periodic job to be run to perform MetroCluster diagnostic.

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

```

* Default value: 1
* Max value: 120
* Min value: 0

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Response

```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Error

```

ONTAP Error Response Codes

```
|===
| Error Code | Description

| 2427132
| MetroCluster is not configured on this cluster.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link
```

```
[cols=3*,options=header]
|===
|Name
```

```

|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

//end collapsible .Definitions block
=====

:leveloffset: -1

= Manage MetroCluster DR groups

:leveloffset: +1

[[IDf35b66a7132fd09a387063d3e7f9daa3]]
= Cluster MetroCluster dr-groups endpoint overview

== Overview

You can use this API to retrieve and display relevant information
pertaining to MetroCluster DR groups. The
`/cluster/metrocluster/dr_groups` endpoint returns a list of all the
dr_groups configured in MetroCluster and their status. Each individual
dr_group can be queried individually using the
`/cluster/metrocluster/dr_groups/{dr_group.id}` endpoint.

[[IDe662b8a1a7c2a8e747e5b8bc6c487716]]
= Retrieve all DR groups in a MetroCluster IP configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
```

```
block]#`/cluster/metrocluster/dr-groups`#
```

Introduced In: 9.8

Retrieves all the DR group in the MetroCluster over IP configuration.

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|partner_cluster.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by partner_cluster.name
```

```
|partner_cluster.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by partner_cluster.uuid
```

```
|id
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by id
```

```
|dr_pairs.partner.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by dr_pairs.partner.name
```

```
|dr_pairs.node.name
```

```
|string
|query
|False
a|Filter by dr_pairs.node.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
```

```
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of Records

|records
|array[link:#metrocluster_dr_group[metrocluster_dr_group]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
```



```

{
  "dr_pairs": [
    {
      "node": {
        "name": "nodeA"
      },
      "partner": {
        "name": "nodeB"
      }
    }
  ],
  "id": 0,
  "partner_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
  }
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

2425734	
---------	--

An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support.

2427132	
---------	--

MetroCluster is not configured on this cluster.

|===

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

Local node of the DR Group.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the node.

```

```
|===
```

```
[#partner]  
[.api-collapsible-fifth-title]  
partner
```

Partner node of the DR Group.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|Name of the node.
```

```
|===
```

```
[#dr_pair]  
[.api-collapsible-fifth-title]  
dr_pair
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|node  
|link:#node[node]  
a|Local node of the DR Group.
```

```
|partner  
|link:#partner[partner]  
a|Partner node of the DR Group.
```

```
|===
```

```

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#partner_cluster]
[.api-collapsible-fifth-title]
partner_cluster

Partner cluster information.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#metrocluster_dr_group]

```

```

[.api-collapsible-fifth-title]
metrocluster_dr_group

DR group information.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|dr_pairs
|array[link:#dr_pair[dr_pair]]
a|

|id
|integer
a|DR Group ID

|partner_cluster
|link:#partner_cluster[partner_cluster]
a|Partner cluster information.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID1fa2800a6f618fae31e82b5295867956]]
= Create a new DR group in a MetroCluster IP configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
```

```
block]#`/cluster/metrocluster/dr-groups`#
```

Introduced In: 9.8

Creates a new DR group in the MetroCluster over IP configuration.

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|return_timeout
```

```
|integer
```

```
|query
```

```
|False
```

a|The number of seconds to allow the call to execute before returning.

When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

```
|return_records
```

```
|boolean
```

```
|query
```

```
|False
```

a|The default is false. If set to true, the records are returned.

* Default value:

```
|===
```

== Request Body


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|dr_pairs
|array[link:#dr_pair[dr_pair]]
a|

|id
|integer
a|DR Group ID

|partner_cluster
|link:#partner_cluster[partner_cluster]
a|Partner cluster information.

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "dr_pairs": [
    {
      "node": {
        "name": "nodeA"
      },
      "partner": {
        "name": "nodeB"
      }
    }
  ],
  "id": 0,
  "partner_cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}

```

```
====
```

```
== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description

| 2425734
| An internal error occurred. Wait a few minutes, and try the operation
| again. For further assistance, contact technical support.
|===
```

```
== Definitions
```

```
[.api-def-first-level]  
.See Definitions  
[%collapsible%closed]  
//Start collapsible Definitions block
```

```
====
```

```
[#node]  
[.api-collapsible-fifth-title]  
node
```

Local node of the DR Group.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|name  
|string  
a|Name of the node.
```

```
|===
```

```
[#partner]  
[.api-collapsible-fifth-title]  
partner
```

Partner node of the DR Group.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|name  
|string  
a|Name of the node.
```

```
|===
```

```

[#dr_pair]
[.api-collapsible-fifth-title]
dr_pair

[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|Local node of the DR Group.

|partner
|link:#partner[partner]
a|Partner node of the DR Group.

|===

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#partner_cluster]
[.api-collapsible-fifth-title]

```

partner_cluster

Partner cluster information.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|

|uuid

|string

a|

|===

[#metrocluster_dr_group]

[.api-collapsible-fifth-title]

metrocluster_dr_group

DR group information.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|dr_pairs

|array[link:#dr_pair[dr_pair]]

a|

|id

|integer

a|DR Group ID

|partner_cluster

|link:#partner_cluster[partner_cluster]

a|Partner cluster information.

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]

|===

|Name
|Type
|Description

|uuid

|string

a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDda697fd119e87b7a6a117b4a6b7ebd7b]]
= Remove a DR group from a MetroCluster IP configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/cluster/metrocluster/dr-groups/{id}`#

*Introduced In:* 9.8

Remove the DR group from the current MetroCluster over IP configuration

```

specified by the DR group id.

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|id

|string

|path

|True

a|

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning.

When doing a POST, PATCH, or DELETE operation on a single record, the

default is 0 seconds. This means that if an asynchronous operation is

started, the server immediately returns HTTP code 202 (Accepted) along

with a link to the job. If a non-zero value is specified for POST, PATCH,

or DELETE operations, ONTAP waits that length of time to see if the job

completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|
```

```

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

```

```
//end collapsible .Definitions block
====
```

```
[[ID0768c789170fb5d5bd5575d5a5236da9]]
= Retrieve DR group information using the DR group ID

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/cluster/metrocluster/dr-groups/{id}`#
```

Introduced In: 9.8

Retrieves the DR group information specified by the DR group id.

== Parameters

```
[cols=5*,options=header]
|==
```

```
|Name
|Type
|In
|Required
|Description
```

```
|id
|string
|path
|True
a|

|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|==
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|dr_pairs
|array[link:#dr_pair[dr_pair]]
a|

|id
|integer
a|DR Group ID

|partner_cluster
|link:#partner_cluster[partner_cluster]
a|Partner cluster information.
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "dr_pairs": [
    {
      "node": {
        "name": "nodeA"
      },
      "partner": {
        "name": "nodeB"
      }
    }
  ],
  "id": 0,
  "partner_cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
```

```

    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]

```

```
//Start collapsible Definitions block
```

```
====
```

```
[#node]
```

```
[.api-collapsible-fifth-title]
```

```
node
```

Local node of the DR Group.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|Name of the node.
```

```
|===
```

```
[#partner]
```

```
[.api-collapsible-fifth-title]
```

```
partner
```

Partner node of the DR Group.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|Name of the node.
```

```
|===
```

```
[#dr_pair]
```

```
[.api-collapsible-fifth-title]
```

```
dr_pair
```



```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|Local node of the DR Group.
```

```
|partner
|link:#partner[partner]
a|Partner node of the DR Group.
```

```
|===
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```

|self
|link:#href[href]
a|

|===

[#partner_cluster]
[.api-collapsible-fifth-title]
partner_cluster

Partner cluster information.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string

```

```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
:leveloffset: -1
```

= View and update MetroCluster interconnects

```
:leveloffset: +1
```

```
[[ID5c26c93e1ca561b8ae8e63de396d1479]]
```

= Cluster MetroCluster interconnects endpoint overview

== Overview

You can use this API to retrieve and display relevant information pertaining to MetroCluster interconnect status. The ``/cluster/metrocluster/interconnects`` endpoint returns a list of all the interconnects in MetroCluster and their status. Each individual interconnect can be queried individually using the ``+/cluster/metrocluster/interconnects/{node.uuid}/{partner_type}/{adapter}+`` endpoint.

```
'''
```

=== Examples

=== Retrieving MetroCluster interconnects

```
----
```

GET `https://<mgmt-ip>/api/cluster/metrocluster/interconnects`

```
{
  "records": [
    {
      "node": {
        "name": "cluster1_01",
        "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-005056826931"
          }
        }
      },
      "partner_type": "ha",
```

```

        "adapter": "e0f",
        "_links": {
            "self": {
                "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/ha/e0f"
            }
        },
        {
            "node": {
                "name": "cluster1_01",
                "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
                "_links": {
                    "self": {
                        "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
                    }
                },
                "partner_type": "ha",
                "adapter": "e0g",
                "_links": {
                    "self": {
                        "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/ha/e0g"
                    }
                }
            },
            {
                "node": {
                    "name": "cluster1_01",
                    "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
                    "_links": {
                        "self": {
                            "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
                        }
                    },
                    "partner_type": "dr",
                    "adapter": "e0f",
                    "_links": {
                        "self": {
                            "href":

```

```

"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/dr/e0f"
    }
  },
  {
    "node": {
      "name": "cluster1_01",
      "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
        }
      },
      "partner_type": "dr",
      "adapter": "e0g",
      "_links": {
        "self": {
          "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/dr/e0g"
        }
      }
    },
    {
      "node": {
        "name": "cluster1_01",
        "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-
005056826931"
          }
        }
      },
      "partner_type": "aux",
      "adapter": "e0f",
      "_links": {
        "self": {
          "href":
"/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-
005056826931/aux/e0f"
        }
      }
    }
  }
}

```

```

    },
    {
      "node": {
        "name": "cluster1_01",
        "uuid": "6fead8fe-8d81-11e9-b5a9-005056826931",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/6fead8fe-8d81-11e9-b5a9-005056826931"
          }
        }
      },
      "partner_type": "aux",
      "adapter": "e0g",
      "_links": {
        "self": {
          "href": "/api/cluster/metrocluster/interconnects/6fead8fe-8d81-11e9-b5a9-005056826931/aux/e0g"
        }
      }
    },
    {
      "node": {
        "name": "cluster1_02",
        "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-00505682dc8b"
          }
        }
      },
      "partner_type": "ha",
      "adapter": "e0f",
      "_links": {
        "self": {
          "href": "/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-00505682dc8b/ha/e0f"
        }
      }
    },
    {
      "node": {
        "name": "cluster1_02",

```

```

        "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
        "_links": {
            "self": {
                "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
            }
        },
        "partner_type": "ha",
        "adapter": "e0g",
        "_links": {
            "self": {
                "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/ha/e0g"
            }
        }
    },
    {
        "node": {
            "name": "cluster1_02",
            "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
            "_links": {
                "self": {
                    "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
                }
            }
        },
        "partner_type": "dr",
        "adapter": "e0f",
        "_links": {
            "self": {
                "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/dr/e0f"
            }
        }
    },
    {
        "node": {
            "name": "cluster1_02",
            "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
            "_links": {
                "self": {
                    "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-

```



```

00505682dc8b"
        }
    },
    "partner_type": "dr",
    "adapter": "e0g",
    "_links": {
        "self": {
            "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/dr/e0g"
        }
    }
},
{
    "node": {
        "name": "cluster1_02",
        "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
        "_links": {
            "self": {
                "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
            }
        }
    },
    "partner_type": "aux",
    "adapter": "e0f",
    "_links": {
        "self": {
            "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/aux/e0f"
        }
    }
},
{
    "node": {
        "name": "cluster1_02",
        "uuid": "f5435191-8d81-11e9-9d4b-00505682dc8b",
        "_links": {
            "self": {
                "href": "/api/cluster/nodes/f5435191-8d81-11e9-9d4b-
00505682dc8b"
            }
        }
    }
},

```

```

        "partner_type": "aux",
        "adapter": "e0g",
        "_links": {
            "self": {
                "href":
"/api/cluster/metrocluster/interconnects/f5435191-8d81-11e9-9d4b-
00505682dc8b/aux/e0g"
            }
        }
    },
    "num_records": 12,
    "_links": {
        "self": {
            "href": "/api/cluster/metrocluster/interconnects"
        }
    }
}

```

=== Retrieves information about a specific MetroCluster interconnect

<https://<mgmt-ip>/api/cluster/metrocluster/interconnects/774b4fbc-86f9-11e9-9051-005056825c71/aux/e0f>

```

{
    "node": {
        "name": "cluster1_01",
        "uuid": "46147363-9857-11e9-9a55-005056828eb9",
        "_links": {
            "self": {
                "href": "/api/cluster/nodes/46147363-9857-11e9-9a55-
005056828eb9"
            }
        }
    },
    "partner_type": "ha",
    "adapter": "e0f",
    "state": "up",
    "type": "iwarp",
    "_links": {
        "self": {
            "href": "/api/cluster/metrocluster/interconnects/46147363-9857-
11e9-9a55-005056828eb9/ha/e0f"
        }
    }
}

```

```

}
----

[[IDd2fdc327299d3bb148bb3f3da055d0f5]]
= Retrieve interconnect adapter information for nodes in MetroCluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/cluster/metrocluster/interconnects`#

*Introduced In:* 9.8

Retrieves a list of interconnect adapter information for nodes in the
MetroCluster.

== Related ONTAP Commands

* `metrocluster interconnect show`

== Learn more

*
xref:{relative_path}cluster_metrocluster_interconnects_endpoint_overview.h
tml[DOC /cluster/metrocluster/interconnects]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|node.name
|string
|query
|False
a|Filter by node.name

```

```
|node.uuid
|string
|query
|False
a|Filter by node.uuid
```

```
|partner_type
|string
|query
|False
a|Filter by partner_type
```

```
|adapter
|string
|query
|False
a|Filter by adapter
```

```
|state
|string
|query
|False
a|Filter by state
```

```
|type
|string
|query
|False
a|Filter by type
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```

|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of Records

|records
|array[link:#metrocluster_interconnect[metrocluster_interconnect]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "adapter": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "partner_type": "string",

```

```

        "state": "string",
        "type": "string"
    }
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
|
| 2425734
| An internal error occurred. Wait a few minutes, and try the operation
| again. For further assistance, contact technical support.
|
| 2427132
| MetroCluster is not configured on this cluster.
|===

```

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

```

```

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [

```

```

    {
      "code": "string",
      "message": "string"
    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

|===

```

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]

```

```

node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#metrocluster_interconnect]
[.api-collapsible-fifth-title]
metrocluster_interconnect

Data for a MetroCluster interconnect. REST:
/api/cluster/metrocluster/interconnects

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|adapter
|string
a|Adapter

|node

```

```
|link:#node[node]
```

```
a|
```

```
|partner_type
```

```
|string
```

```
a|Partner type
```

```
|state
```

```
|string
```

```
a|Adapter status
```

```
|type
```

```
|string
```

```
a|Adapter type
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

```
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDae5b7103ac9a176296bd1819b506882a]]
= Retrieve information about a MetroCluster interconnect for a partner
type and adapter

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/cluster/metrocluster/interconnects/{node.uuid}/{partner_type}/{ad
apter}`#

*Introduced In:* 9.8

Retrieves information about a MetroCluster Interconnect for a specific
partner type and adapter.

```

== Related ONTAP Commands

* `metrocluster interconnect show`

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|node.uuid

|string

|path

|True

a|Node UUID

|partner_type

|string

|path

|True

a|DR Partner type

|adapter

|string

|path

|True

a|Interconnect adapter.

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|==

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|adapter
|string
a|Adapter

|node
|link:#node[node]
a|

|partner_type
|string
a|Partner type

|state
|string
a|Adapter status

|type
|string
a|Adapter type

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
```

```

    "href": "/api/resourcelink"
  }
},
"adapter": "string",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"partner_type": "string",
"state": "string",
"type": "string"
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 2425734
| An internal error occurred. Wait a few minutes, and try the operation
again. For further assistance, contact technical support.

```

```

| 2427132
| MetroCluster is not configured on this cluster.
|===

```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]

```

```

a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```



```

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]

```

```

a|

|name
|string
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments

```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Retrieve MetroCluster node configurations
```

```
:leveloffset: +1
```

```
[[IDaa8ad5b8d3f53ddf7d05d4536629675e]]
```

```
= Cluster MetroCluster nodes endpoint overview
```

```
== Overview
```

```
Retrieves the configuration information for the nodes in the MetroCluster configuration.
```

```
'''
```

=== Example

GET https://<mgmt-ip>/api/cluster/metrocluster/nodes

```
{
  "records": [
    {
      "dr_group_id": 1,
      "cluster": {
        "name": "cluster1",
        "uuid": "8f77de32-9857-11e9-9a55-005056828eb9",
        "_links": {
          "self": {
            "href": "/api/cluster/8f77de32-9857-11e9-9a55-005056828eb9"
          }
        }
      },
      "node": {
        "name": "cluster1_01",
        "uuid": "46147363-9857-11e9-9a55-005056828eb9",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/46147363-9857-11e9-9a55-005056828eb9"
          }
        }
      },
      "dr_mirroring_state": "enabled",
      "configuration_state": "configured",
      "_links": {
        "self": {
          "href": "/api/cluster/metrocluster/nodes/46147363-9857-11e9-9a55-005056828eb9"
        }
      }
    },
    {
      "dr_group_id": 1,
      "cluster": {
        "name": "cluster1",
        "uuid": "8f77de32-9857-11e9-9a55-005056828eb9",
        "_links": {
          "self": {
            "href": "/api/cluster/8f77de32-9857-11e9-9a55-005056828eb9"
          }
        }
      }
    }
  ]
}
```

```

    }
  },
  "node": {
    "name": "cluster1_02",
    "uuid": "cf1dc67f-9857-11e9-bf80-005056829db6",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/cf1dc67f-9857-11e9-bf80-005056829db6"
      }
    }
  },
  "dr_mirroring_state": "enabled",
  "configuration_state": "configured",
  "_links": {
    "self": {
      "href": "/api/cluster/metrocluster/nodes/cf1dc67f-9857-11e9-bf80-005056829db6"
    }
  }
},
{
  "dr_group_id": 1,
  "cluster": {
    "name": "cluster3",
    "uuid": "aa8aa15a-9857-11e9-80c9-00505682e684",
    "_links": {
      "self": {
        "href": "/api/cluster/aa8aa15a-9857-11e9-80c9-00505682e684"
      }
    }
  },
  "node": {
    "name": "cluster3_01",
    "uuid": "5b3b983b-9857-11e9-80c9-00505682e684",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/5b3b983b-9857-11e9-80c9-00505682e684"
      }
    }
  },
  "dr_mirroring_state": "enabled",
  "configuration_state": "configured",

```

```

    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/nodes/5b3b983b-9857-11e9-80c9-00505682e684"
      }
    },
    {
      "dr_group_id": 1,
      "cluster": {
        "name": "cluster3",
        "uuid": "aa8aa15a-9857-11e9-80c9-00505682e684",
        "_links": {
          "self": {
            "href": "/api/cluster/aa8aa15a-9857-11e9-80c9-00505682e684"
          }
        }
      },
      "node": {
        "name": "cluster3_02",
        "uuid": "45bff538-9858-11e9-a624-005056820377",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/45bff538-9858-11e9-a624-005056820377"
          }
        }
      },
      "dr_mirroring_state": "enabled",
      "configuration_state": "configured",
      "_links": {
        "self": {
          "href": "/api/cluster/metrocluster/nodes/45bff538-9858-11e9-a624-005056820377"
        }
      }
    }
  ],
  "num_records": 4,
  "_links": {
    "self": {
      "href": "/api/cluster/metrocluster/nodes?fields=%2A"
    }
  }
}

```

'''

[[IDe289ccf49926b2219028a2de56a87aba]]
= Retrieve MetroCluster nodes and configurations

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/cluster/metrocluster/nodes`#

Introduced In: 9.8

Retrieves MetroCluster nodes and their configurations.

== Related ONTAP Commands

* `metrocluster node show`

== Learn more

*

xref:{relative_path}cluster_metrocluster_nodes_endpoint_overview.html[DOC
/cluster/metrocluster/nodes]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|node.name
|string
|query
|False
a|Filter by node.name

```
|node.uuid
|string
|query
|False
a|Filter by node.uuid

|cluster.name
|string
|query
|False
a|Filter by cluster.name

|cluster.uuid
|string
|query
|False
a|Filter by cluster.uuid

|configuration_state
|string
|query
|False
a|Filter by configuration_state

|dr_group_id
|integer
|query
|False
a|Filter by dr_group_id

|dr_mirroring_state
|string
|query
|False
a|Filter by dr_mirroring_state

|fields
|array[string]
|query
|False
a|Specify the fields to return.
```



```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of Records

|records
|array[link:#metrocluster_node[metrocluster_node]]
a|

|===
```

```
.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    }
  ]
}
```

```

    }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "configuration_state": "string",
  "dr_group_id": 0,
  "dr_mirroring_state": "string",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 2425734
| An internal error occurred. Wait a few minutes, and try the operation
again. For further assistance, contact technical support.
|===

```

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]

```

```

a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```
[#collection_links]
[.api-collapsible-fifth-title]
collection_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|next
|link:#href[href]
a|
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#self_link]
[.api-collapsible-fifth-title]
self_link
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
```

```

|Description

|self
|link:#href[href]
a|

|===

[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#metrocluster_node]
[.api-collapsible-fifth-title]
metrocluster_node
```

Data for a node in a MetroCluster. REST: /api/cluster/metrocluster/nodes

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#self_link[self_link]
a|
```

```
|cluster
|link:#cluster[cluster]
a|
```

```
|configuration_state
|string
a|Configuration state of the node.
```

```
|dr_group_id
|integer
a|DR Group ID.
```

```
|dr_mirroring_state
|string
a|State of the DR mirroring configuration.
```

```

|node
|link:#node[node]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code

```



```

|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDb065fce2742e8500be1960d7f30185d2]]
= Retrieve a node configuration in MetroCluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/cluster/metrocluster/nodes/{node.uuid}`#

*Introduced In:* 9.8

Retrieves the node configuration in the MetroCluster.

== Related ONTAP Commands

* `metrocluster node show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

```

```

|node.uuid
|string
|path
|True
a|Node UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|cluster
|link:#cluster[cluster]
a|

|configuration_state
|string
a|Configuration state of the node.

|dr_group_id
|integer
a|DR Group ID.

|dr_mirroring_state
|string

```

a|State of the DR mirroring configuration.

```
|node
|link:#node[node]
a|
```

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "configuration_state": "string",
  "dr_group_id": 0,
  "dr_mirroring_state": "string",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

====

== Error

ONTAP Error Response Codes

|===

| Error Code | Description

| 2425734

| An internal error occurred. Wait a few minutes, and try the operation again. For further assistance, contact technical support.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

=====

```

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links

```

```

|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Retrieve MetroCluster operations

:leveloffset: +1

[[ID4162c80a8f7e9905b1bb7cea3d31c03d]]
= Cluster MetroCluster operations endpoint overview

= Overview

Retrieves a list of recent MetroCluster operations. To view more
information about a specific operation, use the
`+/cluster/metrocluster/operations/{uuid}+` API endpoint.

```


'''

== Examples

=== Retrieves all MetroCluster operations

GET https://<mgmt-ip>/api/cluster/metrocluster/operations?fields=*

```
{
  "records": [
    {
      "uuid": "a14ae39f-8d85-11e9-b4a7-00505682dc8b",
      "type": "check",
      "state": "successful",
      "start_time": "2019-06-14T11:15:00-07:00",
      "end_time": "2019-06-14T11:16:08-07:00",
      "_links": {
        "self": {
          "href": "/api/cluster/metrocluster/operations/a14ae39f-8d85-11e9-b4a7-00505682dc8b"
        }
      }
    },
    {
      "uuid": "7058df27-8d85-11e9-bbc9-005056826931",
      "type": "configure",
      "state": "successful",
      "start_time": "2019-06-12T19:46:27-07:00",
      "end_time": "2019-06-12T19:48:17-07:00",
      "_links": {
        "self": {
          "href": "/api/cluster/metrocluster/operations/7058df27-8d85-11e9-bbc9-005056826931"
        }
      }
    },
    {
      "uuid": "7849515d-8d84-11e9-bbc9-005056826931",
      "type": "connect",
      "state": "successful",
      "start_time": "2019-06-12T19:39:30-07:00",
      "end_time": "2019-06-12T19:42:02-07:00",
      "_links": {
        "self": {
          "href": "/api/cluster/metrocluster/operations/7849515d-8d84-11e9-bbc9-005056826931"
        }
      }
    }
  ]
}
```

```

    }
  },
  {
    "uuid": "331c79ad-8d84-11e9-b4a7-00505682dc8b",
    "type": "interface_create",
    "state": "successful",
    "start_time": "2019-06-12T19:37:35-07:00",
    "end_time": "2019-06-12T19:37:41-07:00",
    "_links": {
      "self": {
        "href": "/api/cluster/metrocluster/operations/331c79ad-
8d84-11e9-b4a7-00505682dc8b"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/cluster/metrocluster/operations?fields=%2A"
  }
}
}
-----

```

=== Retrieves Information about a specific MetroCluster operation

```

-----
GET https://<mgmt-ip>/api/cluster/metrocluster/operations/0db12274-86fd-
11e9-8053-00505682c342
{
  "uuid": "0db12274-86fd-11e9-8053-00505682c342",
  "name": "check",
  "state": "successful",
  "start_time": "2019-06-06T16:15:01-07:00",
  "end_time": "2019-06-06T16:16:05-07:00",
  "_links": {
    "self": {
      "href": "/api/cluster/metrocluster/operations/0db12274-86fd-
11e9-8053-00505682c342"
    }
  }
}
-----

```

```

[[IDbce65d983af5a6f008c390c623aa4aa7]]
= Retrieve MetroCluster operations on the local cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/cluster/metrocluster/operations`#

*Introduced In:* 9.8

Retrieves the list of MetroCluster operations on the local cluster.

== Related ONTAP Commands

* `metrocluster operation history show`

== Learn more

*
xref:{relative_path}cluster_metrocluster_operations_endpoint_overview.html
[DOC /cluster/metrocluster/operations]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|errors
|string
|query
|False
a|Filter by errors

|uuid
|string

```

```
|query
|False
a|Filter by uuid
```

```
|type
|string
|query
|False
a|Filter by type
```

```
|additional_info
|string
|query
|False
a|Filter by additional_info
```

```
|command_line
|string
|query
|False
a|Filter by command_line
```

```
|end_time
|string
|query
|False
a|Filter by end_time
```

```
|state
|string
|query
|False
a|Filter by state
```

```
|start_time
|string
|query
|False
a|Filter by start_time
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of Records

|records
|array[link:#metrocluster_operation[metrocluster_operation]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
  ],
}
```

```

    "additional_info": "MetroCluster switchover with auto heal completed
successfully.",
    "command_line": "metrocluster switchover",
    "end_time": "2016-03-10T14:35:16-08:00",
    "errors": [
        "siteB (warning): Unable to prepare the partner cluster for a
pending switchback operation. Reason: entry doesn't exist. Reboot the
nodes in the partner cluster before using the \"metrocluster switchback\"
command."
    ],
    "start_time": "2016-03-10T14:33:16-08:00",
    "state": "completed_with_warnings",
    "type": "switchover",
    "uuid": "11111111-2222-3333-4444-abcdefabcdef"
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
|
| 2425734
| An internal error occurred. Wait a few minutes, and try the operation
again. For further assistance, contact technical support.
|===

```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```
|===
```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#collection_links]

```



```

[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#metrocluster_operation]
[.api-collapsible-fifth-title]
metrocluster_operation

Data for a MetroCluster operation. REST:
/api/cluster/metrocluster/operations

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|additional_info
|string
a|Additional information for the auto heal.

|command_line
|string
a|Command line executed with the options specified.

|end_time
|string
a|End Time

|errors
|array[string]
a|List of errors in the operation.

|start_time
|string
a|Start Time

|state
|string
a|Indicates the state of the operation.

|type
|string
a|Name of the operation.

|uuid
|string
a|Identifier for the operation.

```

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

|===

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID1c41f506077f6e921033355481eda671]]
= Retrieve information about a MetroCluster operation

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/cluster/metrocluster/operations/{uuid}`#

*Introduced In:* 9.8

Retrieves information about a specific MetroCluster operation.

== Related ONTAP Commands

* `metrocluster operation show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path

```

```

|True
a|Unique identifier for the operation.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|additional_info
|string
a|Additional information for the auto heal.

|command_line
|string
a|Command line executed with the options specified.

|end_time
|string
a|End Time

|errors
|array[string]
a|List of errors in the operation.

```

```
|start_time
|string
a|Start Time
```

```
|state
|string
a|Indicates the state of the operation.
```

```
|type
|string
a|Name of the operation.
```

```
|uuid
|string
a|Identifier for the operation.
```

```
|===
```

```
.Example response
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "additional_info": "MetroCluster switchover with auto heal completed successfully.",
  "command_line": "metrocluster switchover",
  "end_time": "2016-03-10T14:35:16-08:00",
  "errors": [
    "siteB (warning): Unable to prepare the partner cluster for a pending switchback operation. Reason: entry doesn't exist. Reboot the nodes in the partner cluster before using the \"metrocluster switchback\" command."
  ],
  "start_time": "2016-03-10T14:33:16-08:00",
  "state": "completed_with_warnings",
  "type": "switchover",
  "uuid": "11111111-2222-3333-4444-abcdefabcdef"
}
```

```
====
```

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 2425734
```

```
| An internal error occurred. Wait a few minutes, and try the operation  
again. For further assistance, contact technical support.
```

```
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "error": {
```

```
    "arguments": [
```

```
      {
```

```
        "code": "string",
```

```
        "message": "string"
```

```
      }
```

```
    ],
```

```
    "code": "4",
```

```
    "message": "entry doesn't exist",
```

```
    "target": "uuid"
```

```

    }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]

```


error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
= Manage cluster nodes
```

```
:leveloffset: +1
```

```
[[ID1ba39b13b31572bac2d3f3fe09fcc1da]]
= Cluster nodes endpoint overview
```

```
= Overview
```

You can use this API to add nodes to a cluster, update node-specific configurations, and retrieve the current node configuration details.

```
== Adding a node to a cluster
```

You can add a node to a cluster by issuing a POST `/cluster/nodes` request to a node currently in the cluster. All nodes must be running the same version of ONTAP to use this API. Mixed version joins are not supported in this release. You can provide properties as fields in the body of the POST request to configure node-specific settings. On a successful request, POST `/cluster/nodes` returns a status code of 202 and job information in the body of the request. You can use the `/cluster/jobs` APIs to track the status of the node add job.

```
=== Fields used for adding a node
```

Fields used for the `/cluster/nodes` APIs fall into the following categories:

* Required node fields

- * Optional fields
- * Network interface fields
- * Records field

=== Required node fields

The following field is required for any POST `/cluster/nodes` request:

- * `cluster_interface.ip.address`

=== Optional fields

All of the following fields are used to set up additional cluster-wide configurations:

- * `name`
- * `location`
- * `records`

=== Network interface fields

You can set a node-specific configuration for each node by using the POST `/cluster/nodes` API. If you provide a field in the body of a node, provide it for all nodes in the POST body.

You can provide the node management interface for each node if all node management interfaces in the cluster use the same subnet mask. If the node management interfaces use different subnet masks, use the `/network/ip/interfaces` API to configure the node management interfaces.

=== The records field

To add multiple nodes to the cluster in one request, provide an array named "records" with multiple node entries. Each node entry in "records" must follow the required and optional fields listed previously. When only adding a single node, you do not need a "records" field. See "Examples" for an example of how to use the "records" field.

=== Create recommended aggregates parameter

When you set the "create_recommended_aggregates" parameter to "true", aggregates based on an optimal layout recommended by the system are created on each of the nodes being added to the cluster. The default setting is "false".

'''

== Modifying node configurations

The following fields can be used to modify a node configuration:

- * name
- * location

'''

== Modifying service processor configurations

When modifying the "service_processor" properties, the job returns success immediately if valid network information is passed in. The values remain in their old state until the network information changes have taken effect on the service processor. You can poll the modified properties until the values are updated.

'''

== Deleting a node from a cluster

You can delete a node from the cluster. Before deleting a node from the cluster, shut down all of the node's shared resources, such as virtual interfaces to clients. If any of the node's shared resources are still active, the command fails.

You can use the "force" flag to forcibly remove a node that is down and cannot be brought online to remove its shared resources. This flag is set to "false" by default.

'''

== Node state

The node "state" field in the /cluster/nodes API represents the current operational state of individual nodes.

Note that the state of a node is a transient value and can change depending on the current condition of the node, especially during reboot, takeover, and giveback.

Possible values for the node state are:

- * `_up_` - Node is fully operational and is able to accept and handle management requests. It is connected to a majority of healthy (up) nodes in the cluster through the cluster interconnect and all critical services are online.

- * `_booting_` - Node is starting up and is not yet fully functional. It might not yet be accessible through the management interface or cluster interconnect. One or more critical services are offline on the node and the node is not taken over. The HA partner reports the node's firmware

state as "SF_BOOTING", "SF_BOOTED", or "SF_CLUSTERWAIT".

- * `_down_` - Node is known to be down. It cannot be reached through the management interface or cluster interconnect. The HA partner can be reached and reports that the node is halted/rebooted without takeover. Or, the HA partner cannot be reached (or no SFO configured) but the node shutdown request has been recorded by the quorum change coordinator. The state is reported by the node's HA partner.
- * `_taken_over_` - Node is taken over by its HA partner. The state is reported by the node's HA partner.
- * `_waiting_for_giveback_` - Node is taken over by its HA partner and is now ready and waiting for giveback. To bring the node up, either issue the "giveback" command to the HA partner node or wait for auto-giveback, if enabled. The state is reported by the node's HA partner.
- * `_degraded_` - Node is known to be up but is not yet fully functional. The node can be reached through the cluster interconnect but one or more critical services are offline. Or, the node is not reachable but the node's HA partner can be reached and reports that the node is up with firmware state "SF_UP".
- * `_unknown_` - Node state cannot be determined.

'''

== HA

The "ha" field in the /cluster/nodes API shows the takeover and giveback states of the node along with the current values of the HA fields "enabled" and "auto_giveback".

You can modify the HA fields "enabled" and "auto_giveback", which will change the HA states of the node.

== Performance monitoring

Performance of a node can be monitored by observing the ``metric.++`` and ``statistics.++`` properties. These properties show the performance of a node in terms of cpu utilization. The ``metric.++`` properties denote an average whereas ``statistics.++`` properties denote a real-time monotonically increasing value aggregated across all nodes.

=== Takeover

The takeover "state" field shows the different takeover states of the node. When the state is "failed", the "code" and "message" fields display. Possible values for takeover states are:

- * `_not_attempted_` - Takeover operation is not started and takeover is possible.

```
* _not_possible_ - Takeover operation is not possible. Check the failure
message.
* _in_progress_ - Takeover operation is in progress. The node is taking
over its partner.
* _in_takeover_ - Takeover operation is complete.
* _failed_ - Takeover operation failed. Check the failure message.
```

Possible values for takeover failure code and messages are:

```
* _code_: 852130
_message_: Failed to initiate takeover. Run the "storage failover show-
takeover" command for more information.
* _code_: 852131
_message_: Takeover cannot be completed. Reason: disabled.
```

=== Giveback

The giveback "state" field shows the different giveback states of the node. When the state is "failed", the "code" and "message" fields display. Possible values for giveback states are:

```
* _nothing_to_giveback_ - Node does not have partner aggregates to
giveback.
* _not_attempted_ - Giveback operation is not started.
* _in_progress_ - Giveback operation is in progress.
* _failed_ - Giveback operation failed. Check the failure message.
```

Possible values for giveback failure codes and messages are:

```
* _code_: 852126
_message_: Failed to initiate giveback. Run the "storage failover show-
giveback" command for more information.
```

'''

== Examples

The following examples show how to add nodes to a cluster, update node properties, shutdown and reboot a node, and remove a node from the cluster.

=== Adding a single node with a minimal configuration

```
# Body
add_single_node.txt(body):
```

```

{
  "cluster_interface": {
    "ip": {
      "address": "1.1.1.1"
    }
  }
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/nodes" -d
"@add_single_node.txt"
----

'''

=== Adding multiple nodes in the same request and creating recommended
aggregates

----

# Body
add_multiple_nodes.txt (body):
{
  "records": [
    {
      "name": "node1",
      "cluster_interface": {
        "ip": {
          "address": "1.1.1.1"
        }
      }
    },
    {
      "name": "node2",
      "cluster_interface": {
        "ip": {
          "address": "2.2.2.2"
        }
      }
    }
  ]
}

# Request
curl -X POST "https://<mgmt-
ip>/api/cluster/nodes?create_recommended_aggregates=true" -d

```

```

"@add_multiple_nodes.txt"
----

'''

=== Modifying a cluster-wide configuration

----

# Body
modify_name_and_location.txt(body):
{
"name": "renamedNode",
"location": "newLocation"
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster/nodes" -d
"@modify_name_and_location.txt"
----

'''

=== Shutting down a node

----
curl -X PATCH "https://<mgmt-ip>/api/cluster/nodes/{uuid}?action=shutdown"
----

'''

=== Deleting a node from a cluster

----
curl -X DELETE "https://<mgmt-ip>/api/cluster/nodes/{uuid}"
----

=== Force a node deletion from a cluster

----
curl -X DELETE "https://<mgmt-ip>/api/cluster/nodes/{uuid}?force=true"
----

'''

=== Retrieving the state of all nodes in a cluster

```



```

-----

#Request
curl -siku admin -X GET "https://<mgmt-ip>/api/cluster/nodes?fields=state"

#Response
{
  "records": [
    {
      "uuid": "54440ec3-6127-11e9-a959-005056bb76f9",
      "name": "node2",
      "state": "up",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/54440ec3-6127-11e9-a959-005056bb76f9"
        }
      }
    },
    {
      "uuid": "e02dbef1-6126-11e9-b8fb-005056bb9ce4",
      "name": "node1",
      "state": "up",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e02dbef1-6126-11e9-b8fb-005056bb9ce4"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/cluster/nodes?fields=state"
    }
  }
}
-----

'''

=== Retrieving statistics and metric for a node

In this example, the API returns the "statistics" and "metric" properties.

-----

```

```

#Request
curl -siku admin -X GET "https://<mgmt-
ip>/api/cluster/nodes?fields=statistics,metric"

#Response
{
  "records": [
    {
      "uuid": "6b29327b-21ca-11ea-99aa-005056bb420b",
      "name": "prij-vsml",
      "metric": {
        "timestamp": "2019-12-19T15:50:45Z",
        "duration": "PT15S",
        "status": "ok",
        "processor_utilization": 3
      },
      "statistics": {
        "timestamp": "2019-12-19T15:50:48Z",
        "status": "ok",
        "processor_utilization_raw": 6409411622,
        "processor_utilization_base": 74330229886
      }
    },
    {
      "num_records": 1
    }
  ],
  "num_records": 1
}
----

'''

=== Retrieving takeover and giveback failure codes and messages

```

```
#Request<br> curl -siku admin -X GET "https://<mgmt-ip>/api/cluster/nodes?fields=ha"</mgmt-ip>
```

```

#Response
{
  "records": [
    {
      "uuid": "54440ec3-6127-11e9-a959-005056bb76f9",
      "name": "node2",
      "ha": {
        "enabled": false,
        "auto_giveback": false,
        "partners": [
          {
            "uuid": "e02dbef1-6126-11e9-b8fb-005056bb9ce4",
            "name": "node1"
          }
        ]
      }
    }
  ]
}

```

```

],
"giveback": {
"state": "nothing_to_giveback"
},
"takeover": {
"state": "not_possible",
"failure": {
"message": "Takeover cannot be completed. Reason: disabled.",
"code": 852131
}
},
"ports": [
{
"name": "e0h"
},
{
"name": "N/A"
}
],
"_links": {
"self": {
"href": "/api/cluster/nodes/54440ec3-6127-11e9-a959-005056bb76f9"
}
},
{
"uuid": "e02dbef1-6126-11e9-b8fb-005056bb9ce4",
"name": "node1",
"ha": {
"enabled": false,
"auto_giveback": false,
"partners": [
{
"uuid": "54440ec3-6127-11e9-a959-005056bb76f9",
"name": "node2"
}
],
"giveback": {
"state": "nothing_to_giveback"
},
"takeover": {
"state": "not_possible",
"failure": {
"message": "Takeover cannot be completed. Reason: disabled.",
"code": 852131
}
},
"ports": [
{
"name": "e0h"
},
{
"name": "N/A"
}
]
}

```

```

}
],
},
"_links": {
  "self": {
    "href": "/api/cluster/nodes/e02dbef1-6126-11e9-b8fb-005056bb9ce4"
  }
},
},
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/cluster/nodes?fields=state"
  }
},
}
}
}

```

Retrieve nodes in a cluster

GET `/cluster/nodes`

Introduced In: 9.6

Retrieves the nodes in the cluster.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `statistics.*`
- `metric.*`

Related ONTAP commands

- `system node show`

Parameters

Name	Type	In	Required	Description
<code>service_processor.link_status</code>	string	query	False	Filter by <code>service_processor.link_status</code>
<code>service_processor.state</code>	string	query	False	Filter by <code>service_processor.state</code>

Name	Type	In	Required	Description
service_processor.firmware_version	string	query	False	Filter by service_processor.firmware_version
service_processor.mac_address	string	query	False	Filter by service_processor.mac_address
service_processor.dhcp_enabled	boolean	query	False	Filter by service_processor.dhcp_enabled
service_processor.ipv4_interface.address	string	query	False	Filter by service_processor.ipv4_interface.address
service_processor.ipv4_interface.gateway	string	query	False	Filter by service_processor.ipv4_interface.gateway
service_processor.ipv4_interface.netmask	string	query	False	Filter by service_processor.ipv4_interface.netmask
service_processor.ipv6_interface.address	string	query	False	Filter by service_processor.ipv6_interface.address
service_processor.ipv6_interface.gateway	string	query	False	Filter by service_processor.ipv6_interface.gateway
service_processor.ipv6_interface.netmask	string	query	False	Filter by service_processor.ipv6_interface.netmask
serial_number	string	query	False	Filter by serial_number

Name	Type	In	Required	Description
statistics.timestamp	string	query	False	Filter by statistics.timestamp <ul style="list-style-type: none"> • Introduced in: 9.8
statistics.processor_utilization_base	integer	query	False	Filter by statistics.processor_utilization_base <ul style="list-style-type: none"> • Introduced in: 9.8
statistics.processor_utilization_raw	integer	query	False	Filter by statistics.processor_utilization_raw <ul style="list-style-type: none"> • Introduced in: 9.8
statistics.status	string	query	False	Filter by statistics.status <ul style="list-style-type: none"> • Introduced in: 9.8
state	string	query	False	Filter by state <ul style="list-style-type: none"> • Introduced in: 9.7
vendor_serial_number	string	query	False	Filter by vendor_serial_number <ul style="list-style-type: none"> • Introduced in: 9.7
location	string	query	False	Filter by location
vm.provider_type	string	query	False	Filter by vm.provider_type <ul style="list-style-type: none"> • Introduced in: 9.7
uptime	integer	query	False	Filter by uptime

Name	Type	In	Required	Description
metrocluster.type	string	query	False	Filter by metrocluster.type <ul style="list-style-type: none"> • Introduced in: 9.8
cluster_interfaces.name	string	query	False	Filter by cluster_interfaces.name
cluster_interfaces.uuid	string	query	False	Filter by cluster_interfaces.uuid
cluster_interfaces.ip.address	string	query	False	Filter by cluster_interfaces.ip.address
system_machine_type	string	query	False	Filter by system_machine_type <ul style="list-style-type: none"> • Introduced in: 9.7
management_interfaces.name	string	query	False	Filter by management_interfaces.name
management_interfaces.uuid	string	query	False	Filter by management_interfaces.uuid
management_interfaces.ip.address	string	query	False	Filter by management_interfaces.ip.address
system_id	string	query	False	Filter by system_id <ul style="list-style-type: none"> • Introduced in: 9.7
controller.frus.state	string	query	False	Filter by controller.frus.state
controller.frus.id	integer	query	False	Filter by controller.frus.id

Name	Type	In	Required	Description
controller.frus.type	string	query	False	Filter by controller.frus.type
controller.over_temperature	string	query	False	Filter by controller.over_temperature
controller.flash_cache.serial_number	string	query	False	Filter by controller.flash_cache.serial_number
controller.flash_cache.slot	string	query	False	Filter by controller.flash_cache.slot
controller.flash_cache.state	string	query	False	Filter by controller.flash_cache.state
controller.flash_cache.hardware_revision	string	query	False	Filter by controller.flash_cache.hardware_revision
controller.flash_cache.part_number	string	query	False	Filter by controller.flash_cache.part_number
controller.flash_cache.capacity	integer	query	False	Filter by controller.flash_cache.capacity
controller.flash_cache.firmware_version	string	query	False	Filter by controller.flash_cache.firmware_version
controller.flash_cache.model	string	query	False	Filter by controller.flash_cache.model
date	string	query	False	Filter by date
uuid	string	query	False	Filter by uuid
ha.partners.name	string	query	False	Filter by ha.partners.name

Name	Type	In	Required	Description
ha.partners.uuid	string	query	False	Filter by ha.partners.uuid
ha.ports.number	integer	query	False	Filter by ha.ports.number • Introduced in: 9.7
ha.ports.state	string	query	False	Filter by ha.ports.state • Introduced in: 9.7
ha.giveback.failure.code	integer	query	False	Filter by ha.giveback.failure.code • Introduced in: 9.7
ha.giveback.failure.message	string	query	False	Filter by ha.giveback.failure.message • Introduced in: 9.7
ha.giveback.state	string	query	False	Filter by ha.giveback.state • Introduced in: 9.7
ha.enabled	boolean	query	False	Filter by ha.enabled
ha.takeover.failure.message	string	query	False	Filter by ha.takeover.failure.message • Introduced in: 9.7

Name	Type	In	Required	Description
ha.takeover.failure.code	integer	query	False	Filter by ha.takeover.failure.code • Introduced in: 9.7
ha.takeover.state	string	query	False	Filter by ha.takeover.state • Introduced in: 9.7
ha.auto_giveback	boolean	query	False	Filter by ha.auto_giveback
metric.duration	string	query	False	Filter by metric.duration • Introduced in: 9.8
metric.processor_utilization	integer	query	False	Filter by metric.processor_utilization • Introduced in: 9.8
metric.timestamp	string	query	False	Filter by metric.timestamp • Introduced in: 9.8
metric.status	string	query	False	Filter by metric.status • Introduced in: 9.8
model	string	query	False	Filter by model
name	string	query	False	Filter by name
version.full	string	query	False	Filter by version.full

Name	Type	In	Required	Description
version.minor	integer	query	False	Filter by version.minor
version.major	integer	query	False	Filter by version.major
version.generation	integer	query	False	Filter by version.generation
membership	string	query	False	Filter by membership
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cluster_interfaces": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "ip": {
            "address": "10.10.10.7"
          },
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "controller": {
        "flash_cache": [
          {
            "capacity": "1024000000000",
            "firmware_version": "NA05",
            "hardware_revision": "A1",
            "model": "X1970A",
            "part_number": "119-00207",
            "serial_number": "A22P5061550000187",
            "slot": "6-1",
            "state": "string"
          }
        ],
        "frus": [
```

```

        {
            "id": 0,
            "state": "string",
            "type": "string"
        }
    ],
    "over_temperature": "string"
},
"date": "2019-04-17T11:49:26-04:00",
"ha": {
    "giveback": {
        "failure": {
            "code": "852126",
            "message": "Failed to initiate giveback. Run the \"storage
failover show-giveback\" command for more information."
        },
        "state": "failed"
    },
    "partners": [
        {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
    ],
    "ports": [
        {
            "number": "0",
            "state": "active"
        }
    ],
    "takeover": {
        "failure": {
            "code": "852130",
            "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
        },
        "state": "failed"
    }
},
"location": "rack 2 row 5",
"management_interfaces": [

```

```

    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "membership": "string",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "processor_utilization": "13",
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "metrocluster": {
    "type": "string"
  },
  "model": "FAS3070",
  "name": "node-01",
  "serial_number": "4048820-60-9",
  "service_processor": {
    "firmware_version": "string",
    "ipv4_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "ipv6_interface": {
      "address": "10.10.10.7",
      "gateway": "10.1.1.1",
      "netmask": "24"
    },
    "link_status": "string",
    "mac_address": "string",
    "state": "string"
  }
}

```

```

    },
    "state": "string",
    "statistics": {
      "processor_utilization_base": "12345123",
      "processor_utilization_raw": "13",
      "status": "ok",
      "timestamp": "2017-01-25T11:20:13Z"
    },
    "system_id": "0537035403",
    "system_machine_type": "7Y56-CTOWW1",
    "uptime": "300536",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
    "vendor_serial_number": "791603000068",
    "version": {
      "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
      "generation": "9",
      "major": "4",
      "minor": "0"
    },
    "vm": {
      "provider_type": "string"
    }
  }
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information

Name	Type	Description
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	List of FRUs on the node. Only returned when requested by name.

Name	Type	Description
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

failure

Indicates the failure code and message.

Name	Type	Description
code	integer	Message code
message	string	Detailed message based on the state.

giveback

Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ports

Name	Type	Description
number	integer	HA port number

Name	Type	Description
state	string	<p>HA port state:</p> <ul style="list-style-type: none"> • <i>down</i> - Logical HA link is down. • <i>initialized</i> - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port. • <i>armed</i> - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port. • <i>active</i> - Logical HA link is active. • <i>reserved</i> - Logical HA link is active, but the physical link is down.

takeover

This represents the state of the node that is taking over storage from its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
giveback	giveback	Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
partners	array[partners]	Nodes in this node's High Availability (HA) group.
ports	array[ports]	
takeover	takeover	This represents the state of the node that is taking over storage from its HA partner.

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

metric

CPU performance for the nodes.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
processor_utilization	integer	Average CPU Utilization for the node

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

metrocluster

Metrocluster

Name	Type	Description
type	string	The Metrocluster configuration type

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to "true" to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

statistics

Raw CPU performance for the nodes.

Name	Type	Description
processor_utilization_base	integer	Base counter for CPU Utilization.
processor_utilization_raw	integer	Raw CPU Utilization for the node. This should be divided by the processor_utilization_base to calculate the percentage CPU utilization for the node.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.

Name	Type	Description
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

vm

Name	Type	Description
provider_type	string	Cloud provider where the VM is hosted.

records

Complete node information

Name	Type	Description
_links	_links	
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none"> • example: 2019-04-17T11:49:26-04:00 • format: date-time • readOnly: 1
ha	ha	
location	string	
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
metric	metric	CPU performance for the nodes.
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	

Name	Type	Description
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
statistics	statistics	Raw CPU performance for the nodes.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.
vm	vm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Add a node or nodes to a cluster

POST /cluster/nodes

Introduced In: 9.6

Adds a node or nodes to the cluster.

Required properties

- `cluster_interface.ip.address`

Related ONTAP commands

- `cluster add-node`
- `network interface create`
- `storage aggregate auto-provision`
- `system node modify`
- `system service-processor network modify`

Parameters

Name	Type	In	Required	Description
create_recommended_aggregates	boolean	query	False	<p>Creates aggregates based on an optimal layout recommended by the system.</p> <ul style="list-style-type: none"> • Default value: • Introduced in: 9.7
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none">• example: 2019-04-17T11:49:26-04:00• format: date-time• readOnly: 1• Introduced in: 9.6• x-nullable: true
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	

Name	Type	Description
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.
vm	vm	

Example request

```
{
  "cluster_interface": {
    "ip": {
      "address": "10.10.10.7"
    }
  },
  "cluster_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "controller": {
    "flash_cache": [
      {
        "capacity": "1024000000000",
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
        "slot": "6-1",
        "state": "string"
      }
    ],
    "frus": [
      {
        "id": 0,
        "state": "string",
        "type": "string"
      }
    ],
    "over_temperature": "string"
  },
  "date": "2019-04-17T11:49:26-04:00",
  "ha": {
    "giveback": {
      "failure": {
        "code": "852126",
        "message": "Failed to initiate giveback. Run the \"storage failover show-giveback\" command for more information."
      },
      "state": "failed"
    }
  },
}
```

```

"partners": [
  {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"ports": [
  {
    "number": "0",
    "state": "active"
  }
],
"takeover": {
  "failure": {
    "code": "852130",
    "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
  },
  "state": "failed"
}
},
"location": "rack 2 row 5",
"management_interface": {
  "ip": {
    "address": "10.10.10.7"
  }
},
"management_interfaces": [
  {
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"membership": "string",
"metrocluster": {
  "type": "string"
},
"model": "FAS3070",
"name": "node-01",
"serial_number": "4048820-60-9",
"service_processor": {
  "dhcp_enabled": null,
  "firmware_version": "string",
  "ipv4_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",

```

```

    "netmask": "24"
  },
  "link_status": "string",
  "mac_address": "string",
  "state": "string"
},
"state": "string",
"system_id": "0537035403",
"system_machine_type": "7Y56-CTOWW1",
"uptime": "300536",
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
"vendor_serial_number": "791603000068",
"vm": {
  "provider_type": "string"
}
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

{
  "job": {
    "uuid": "string"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
262245	The value provided was invalid.

Error Code	Description
1179795	A node being added is already in the cluster.
1179813	Fields set for one node must be set for all nodes.
1179817	The IP address, subnet mask, and gateway must all be provided for cluster management interface.
1179818	The IP address and gateway must be of the same family.
1179821	An IP address and subnet mask conflicts with an existing entry.
131727360	A node cannot be added to the cluster. This is a generic code, see response message for details.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

ip

IP information

cluster_interfaces

Network interface

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	

Name	Type	Description
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	List of FRUs on the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

failure

Indicates the failure code and message.

Name	Type	Description
code	integer	Message code

Name	Type	Description
message	string	Detailed message based on the state.

giveback

Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

partners

Name	Type	Description
name	string	
uuid	string	

ports

Name	Type	Description
number	integer	HA port number
state	string	HA port state: <ul style="list-style-type: none"> • <i>down</i> - Logical HA link is down. • <i>initialized</i> - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port. • <i>armed</i> - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port. • <i>active</i> - Logical HA link is active. • <i>reserved</i> - Logical HA link is active, but the physical link is down.

takeover

This represents the state of the node that is taking over storage from its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
giveback	giveback	Represents the state of the node that is giving storage back to its HA partner.
partners	array[partners]	Nodes in this node's High Availability (HA) group.
ports	array[ports]	
takeover	takeover	This represents the state of the node that is taking over storage from its HA partner.

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

Name	Type	Description
ip	node_setup_ip	The IP configuration for cluster setup.

management_interfaces

Network interface

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

metric

CPU performance for the nodes.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
processor_utilization	integer	Average CPU Utilization for the node
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
timestamp	string	The timestamp of the performance data.

metrocluster

Metrocluster

Name	Type	Description
type	string	The Metrocluster configuration type

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

statistics

Raw CPU performance for the nodes.

Name	Type	Description
processor_utilization_base	integer	Base counter for CPU Utilization.
processor_utilization_raw	integer	Raw CPU Utilization for the node. This should be divided by the processor_utilization_base to calculate the percentage CPU utilization for the node.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
timestamp	string	The timestamp of the performance data.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

vm

Name	Type	Description
provider_type	string	Cloud provider where the VM is hosted.

node

Complete node information

Name	Type	Description
cluster_interface	cluster_interface	The cluster network IP address of the node to be added.
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information

Name	Type	Description
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none"> • example: 2019-04-17T11:49:26-04:00 • format: date-time • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
ha	ha	
location	string	
management_interface	management_interface	The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	

Name	Type	Description
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.
vm	vm	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a node from a cluster

DELETE /cluster/nodes/{uuid}

Introduced In: 9.7

Deletes a node from the cluster.

Note that before deleting a node from the cluster, you must shut down all of the node's shared resources, such as virtual interfaces to clients. If any of the node's shared resources are still active, the command fails.

Optional parameters:

- `force` - Forcibly removes a node that is down and cannot be brought online to remove its shared resources. This flag is set to "false" by default.

Related ONTAP commands

- `cluster remove-node`

Learn more

- [DOC /cluster/nodes](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

Name	Type	In	Required	Description
force	boolean	query	False	<p>Set the force flag to "true" to forcibly remove a node that is down and cannot be brought online to remove its shared resources.</p> <ul style="list-style-type: none"> • Default value:
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
458755	Replication service is offline.
458758	Failed to load job for cluster remove node operation as the job exists.
1179732	Cannot remove a node in a single-node cluster.
1179735	Node is not part of a cluster.
1182805	Cannot remove a node from the node network address of the node to be removed.
2293765	Removing a node only works for nodes not in failover configuration.
2293767	Node has volumes. Either move or delete them from the node before removing the node.
2293768	Node is the home node for one or more logical interfaces.
2293769	Node is the current node for one or more logical interfaces.
2293770	Node has data logical interfaces configured as target node.

Error Code	Description
2293789	Removing a node only works for nodes not in HA configuration.
2293796	Cluster ring is offline on the node
2293798	Cannot forcibly remove a node that is online.
2293800	Node is configured with MetroCluster.
2293801	Cannot remove node because it has foreign LUN Imports.
2293812	Node is a member of MetroCluster DR group.
2293813	Cannot remove a node from the cluster because a controller replacement is in progress.
2293814	The DELETE operation is not supported until the cluster is upgraded.
2293816	Cannot remove node because its Storage Encryption devices use authentication keys (AKs) that will not be available to the node after it leaves the cluster.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve node information

GET /cluster/nodes/{uuid}

Introduced In: 9.6

Retrieves information for the node.

Related ONTAP commands

- `cluster add-node-status`
- `cluster date show`
- `cluster ha show`
- `network interface show`
- `network port show`
- `storage failover show`
- `system controller show`
- `system node show`
- `system node show-discovered`
- `system service-processor network show`
- `system service-processor show`
- `version`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none">• format: uuid
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
<code>_links</code>	<code>_links</code>	
<code>cluster_interfaces</code>	<code>array[cluster_interfaces]</code>	
<code>controller</code>	<code>controller</code>	Controller information

Name	Type	Description
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none"> • example: 2019-04-17T11:49:26-04:00 • format: date-time • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
ha	ha	
location	string	
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.

Name	Type	Description
metric	metric	CPU performance for the nodes.
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
statistics	statistics	Raw CPU performance for the nodes.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.

Name	Type	Description
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.
vm	vm	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster_interfaces": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "controller": {
    "flash_cache": [
      {
        "capacity": "1024000000000",
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
        "slot": "6-1",
        "state": "string"
      }
    ],
    "frus": [
      {
        "id": 0,
        "state": "string",
        "type": "string"
      }
    ],
    "over_temperature": "string"
  },
  "date": "2019-04-17T11:49:26-04:00",
  "ha": {
```

```

"giveback": {
  "failure": {
    "code": "852126",
    "message": "Failed to initiate giveback. Run the \"storage
failover show-giveback\" command for more information."
  },
  "state": "failed"
},
"partners": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"ports": [
  {
    "number": "0",
    "state": "active"
  }
],
"takeover": {
  "failure": {
    "code": "852130",
    "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
  },
  "state": "failed"
},
"location": "rack 2 row 5",
"management_interfaces": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",

```

```

    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
],
"membership": "string",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"duration": "PT15S",
"processor_utilization": "13",
"status": "ok",
"timestamp": "2017-01-25T11:20:13Z"
},
"metrocluster": {
  "type": "string"
},
"model": "FAS3070",
"name": "node-01",
"serial_number": "4048820-60-9",
"service_processor": {
  "firmware_version": "string",
  "ipv4_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "ipv6_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
},
"link_status": "string",
"mac_address": "string",
"state": "string"
},
"state": "string",
"statistics": {
  "processor_utilization_base": "12345123",
  "processor_utilization_raw": "13",
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"system_id": "0537035403",
"system_machine_type": "7Y56-CTOWW1",

```

```

"uptime": "300536",
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
"vendor_serial_number": "791603000068",
"version": {
  "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
  "generation": "9",
  "major": "4",
  "minor": "0"
},
"vm": {
  "provider_type": "string"
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

ip

IP information

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	List of FRUs on the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

failure

Indicates the failure code and message.

Name	Type	Description
code	integer	Message code
message	string	Detailed message based on the state.

giveback

Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

partners

Name	Type	Description
_links	_links	
name	string	
uuid	string	

ports

Name	Type	Description
number	integer	HA port number

Name	Type	Description
state	string	<p>HA port state:</p> <ul style="list-style-type: none"> • <i>down</i> - Logical HA link is down. • <i>initialized</i> - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port. • <i>armed</i> - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port. • <i>active</i> - Logical HA link is active. • <i>reserved</i> - Logical HA link is active, but the physical link is down.

takeover

This represents the state of the node that is taking over storage from its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
giveback	giveback	Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
partners	array[partners]	Nodes in this node's High Availability (HA) group.
ports	array[ports]	
takeover	takeover	This represents the state of the node that is taking over storage from its HA partner.

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

Name	Type	Description
_links	_links	
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

metric

CPU performance for the nodes.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
processor_utilization	integer	Average CPU Utilization for the node

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

metrocluster

Metrocluster

Name	Type	Description
type	string	The Metrocluster configuration type

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to "true" to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	
state	string	

statistics

Raw CPU performance for the nodes.

Name	Type	Description
processor_utilization_base	integer	Base counter for CPU Utilization.
processor_utilization_raw	integer	Raw CPU Utilization for the node. This should be divided by the processor_utilization_base to calculate the percentage CPU utilization for the node.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.

Name	Type	Description
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

vm

Name	Type	Description
provider_type	string	Cloud provider where the VM is hosted.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update node information

PATCH /cluster/nodes/{uuid}

Introduced In: 9.6

Updates the node information or performs shutdown/reboot actions on a node.

Related ONTAP commands

- `cluster ha modify`
- `storage failover modify`
- `system node modify`

- system node reboot
- system service-processor network modify

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none"> • format: uuid
action	string	query	False	<p>The shutdown action shuts the node down and transfers storage control to its HA group if storage failover is enabled. The reboot action reboots the node and transfers storage control to its HA group if storage failover is enabled. The giveback action transfers storage control back to the owner from its HA group.</p> <ul style="list-style-type: none"> • enum: ["shutdown", "reboot", "giveback"]
shutdown_reboot_reason	string	query	False	Indicates the reason for the reboot or shutdown. This only applies when an action of reboot or shutdown is provided.
allow_data_outage	boolean	query	False	<p>This only applies when an action of reboot or shutdown is provided. It allows storage failover to be bypassed along with any failures related to maintaining quorum in the cluster.</p> <ul style="list-style-type: none"> • Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information

Name	Type	Description
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none"> • example: 2019-04-17T11:49:26-04:00 • format: date-time • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
ha	ha	
location	string	
management_interfaces	array[management_interfaces]	
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.

Name	Type	Description
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.
vm	vm	

Example request

```
{
  "cluster_interfaces": [
    {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "controller": {
    "flash_cache": [
      {
        "capacity": "1024000000000",
        "firmware_version": "NA05",
        "hardware_revision": "A1",
        "model": "X1970A",
        "part_number": "119-00207",
        "serial_number": "A22P5061550000187",
        "slot": "6-1",
        "state": "string"
      }
    ],
    "frus": [
      {
        "id": 0,
        "state": "string",
        "type": "string"
      }
    ],
    "over_temperature": "string"
  },
  "date": "2019-04-17T11:49:26-04:00",
  "ha": {
    "giveback": {
      "failure": {
        "code": "852126",
        "message": "Failed to initiate giveback. Run the \"storage failover show-giveback\" command for more information."
      },
      "state": "failed"
    },
    "partners": [
      {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  }
}
```

```

],
"ports": [
  {
    "number": "0",
    "state": "active"
  }
],
"takeover": {
  "failure": {
    "code": "852130",
    "message": "Failed to initiate takeover. Run the \"storage
failover show-takeover\" command for more information."
  },
  "state": "failed"
}
},
"location": "rack 2 row 5",
"management_interfaces": [
  {
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"membership": "string",
"metrocluster": {
  "type": "string"
},
"model": "FAS3070",
"name": "node-01",
"serial_number": "4048820-60-9",
"service_processor": {
  "firmware_version": "string",
  "ipv4_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "ipv6_interface": {
    "address": "10.10.10.7",
    "gateway": "10.1.1.1",
    "netmask": "24"
  },
  "link_status": "string",
  "mac_address": "string",
  "state": "string"
},
},

```

```
"state": "string",
"system_id": "0537035403",
"system_machine_type": "7Y56-CTOWW1",
"uptime": "300536",
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412",
"vendor_serial_number": "791603000068",
"vm": {
  "provider_type": "string"
}
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
852046	HA partner node
852115	The reboot/shutdown is prevented because LIFs cannot be moved away from the node
3604514	A reboot or shutdown request is already in progress.

Error Code	Description
3604515	Reboot or shutdown of all nodes results in data service failure and client disruption for the entire cluster. Use "allow-data-outage=true" to bypass this check.
9240606	The reboot/shutdown is prevented due to quorum warnings.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

node_setup_ip

The IP configuration for cluster setup.

Name	Type	Description
address	string	IPv4 or IPv6 address

cluster_interface

The cluster network IP address of the node to be added.

ip

IP information

cluster_interfaces

Network interface

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

flash_cache

Name	Type	Description
capacity	integer	Size in bytes
firmware_version	string	
hardware_revision	string	
model	string	
part_number	string	

Name	Type	Description
serial_number	string	
slot	string	
state	string	

frus

Name	Type	Description
id	integer	
state	string	
type	string	

controller

Controller information

Name	Type	Description
flash_cache	array[flash_cache]	A list of Flash-Cache devices. Only returned when requested by name.
frus	array[frus]	List of FRUs on the node. Only returned when requested by name.
over_temperature	string	Specifies whether the hardware is currently operating outside of its recommended temperature range. The hardware shuts down if the temperature exceeds critical thresholds.

failure

Indicates the failure code and message.

Name	Type	Description
code	integer	Message code
message	string	Detailed message based on the state.

giveback

Represents the state of the node that is giving storage back to its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

partners

Name	Type	Description
name	string	
uuid	string	

ports

Name	Type	Description
number	integer	HA port number
state	string	<p>HA port state:</p> <ul style="list-style-type: none"> • <i>down</i> - Logical HA link is down. • <i>initialized</i> - Logical HA link is initialized. The physical link is up, but the subnet manager hasn't started to configure the port. • <i>armed</i> - Logical HA link is armed. The physical link is up and the subnet manager started but did not yet complete configuring the port. • <i>active</i> - Logical HA link is active. • <i>reserved</i> - Logical HA link is active, but the physical link is down.

takeover

This represents the state of the node that is taking over storage from its HA partner.

Name	Type	Description
failure	failure	Indicates the failure code and message.
state	string	

ha

Name	Type	Description
auto_giveback	boolean	Specifies whether giveback is automatically initiated when the node that owns the storage is ready.
enabled	boolean	Specifies whether or not storage failover is enabled.
giveback	giveback	Represents the state of the node that is giving storage back to its HA partner.
partners	array[partners]	Nodes in this node's High Availability (HA) group.
ports	array[ports]	
takeover	takeover	This represents the state of the node that is taking over storage from its HA partner.

management_interface

The management interface of the node to be added. The subnet mask is set based on the management interface of the cluster or the management interfaces of other nodes.

management_interfaces

Network interface

Name	Type	Description
ip	ip	IP information
name	string	The name of the interface.
uuid	string	The UUID that uniquely identifies the interface.

metric

CPU performance for the nodes.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
processor_utilization	integer	Average CPU Utilization for the node
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

metrocluster

Metrocluster

Name	Type	Description
type	string	The Metrocluster configuration type

ipv4_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

ipv6_interface

Object to setup an interface along with its default router.

Name	Type	Description
address	string	IPv4 or IPv6 address
gateway	string	The IPv4 or IPv6 address of the default router.
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

service_processor

Name	Type	Description
dhcp_enabled	boolean	Set to "true" to use DHCP to configure an IPv4 interface.
firmware_version	string	The version of firmware installed.
ipv4_interface	ipv4_interface	Object to setup an interface along with its default router.
ipv6_interface	ipv6_interface	Object to setup an interface along with its default router.
link_status	string	
mac_address	string	

Name	Type	Description
state	string	

statistics

Raw CPU performance for the nodes.

Name	Type	Description
processor_utilization_base	integer	Base counter for CPU Utilization.
processor_utilization_raw	integer	Raw CPU Utilization for the node. This should be divided by the processor_utilization_base to calculate the percentage CPU utilization for the node.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

vm

Name	Type	Description
provider_type	string	Cloud provider where the VM is hosted.

node

Complete node information

Name	Type	Description
cluster_interfaces	array[cluster_interfaces]	
controller	controller	Controller information
date	string	<p>The current or "wall clock" time of the node in ISO-8601 date, time, and time zone format. The ISO-8601 date and time are localized based on the ONTAP cluster's timezone setting.</p> <ul style="list-style-type: none"> • example: 2019-04-17T11:49:26-04:00 • format: date-time • readOnly: 1 • Introduced in: 9.6 • x-nullable: true
ha	ha	
location	string	
management_interfaces	array[management_interfaces]	

Name	Type	Description
membership	string	<p>Possible values:</p> <ul style="list-style-type: none"> • <i>available</i> - A node is detected on the internal cluster network and can be added to the cluster. Nodes that have a membership of "available" are not returned when a GET request is called when the cluster exists. Provide a query on the "membership" property for <i>available</i> to scan for nodes on the cluster network. Nodes that have a membership of "available" are returned automatically before a cluster is created. • <i>joining</i> - Joining nodes are in the process of being added to the cluster. The node might be progressing through the steps to become a member or might have failed. The job to add the node or create the cluster provides details on the current progress of the node. • <i>member</i> - Nodes that are members have successfully joined the cluster.
metrocluster	metrocluster	Metrocluster
model	string	
name	string	
serial_number	string	
service_processor	service_processor	

Name	Type	Description
state	string	<p>State of the node:</p> <ul style="list-style-type: none"> • <i>up</i> - Node is up and operational. • <i>booting</i> - Node is booting up. • <i>down</i> - Node has stopped or is dumping core. • <i>taken_over</i> - Node has been taken over by its HA partner and is not yet waiting for giveback. • <i>waiting_for_giveback</i> - Node has been taken over by its HA partner and is waiting for the HA partner to giveback disks. • <i>degraded</i> - Node has one or more critical services offline. • <i>unknown</i> - Node or its HA partner cannot be contacted and there is no information on the node's state.
system_id	string	
system_machine_type	string	OEM system machine type.
uptime	integer	The total time, in seconds, that the node has been up.
uuid	string	
vendor_serial_number	string	OEM vendor serial number.
vm	vm	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve node historical performance metrics

GET /cluster/nodes/{uuid}/metrics

Introduced In: 9.8

Retrieves historical performance metrics for a node.

Parameters

Name	Type	In	Required	Description
duration	string	query	False	Filter by duration
processor_utilization	integer	query	False	Filter by processor_utilization
timestamp	string	query	False	Filter by timestamp
status	string	query	False	Filter by status
uuid	string	path	True	Unique identifier of the node.

Name	Type	In	Required	Description
interval	string	query	False	<p>The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y. The period for each time range is as follows:</p> <ul style="list-style-type: none"> • 1h: Metrics over the most recent hour sampled over 15 seconds. • 1d: Metrics over the most recent day sampled over 5 minutes. • 1w: Metrics over the most recent week sampled over 30 minutes. • 1m: Metrics over the most recent month sampled over 2 hours. • 1y: Metrics over the most recent year sampled over a day. • Default value: 1 • enum: ["1h", "1d", "1w", "1m", "1y"]

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "processor_utilization": "13",
      "status": "ok",
      "timestamp": "2017-01-25T11:20:13Z"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

records

CPU performance for the nodes.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
processor_utilization	integer	Average CPU Utilization for the node

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Cluster NTP

Cluster NTP endpoint overview

Overview

ONTAP uses the Network Time Protocol (NTP) for world clock time synchronization of the cluster. Some functional services require the time to be correct to within one second for all the nodes in the cluster.

The success and speed of this synchronization depends on the number, alignment, and consistent network latency of external time servers. It is a best practice to configure ONTAP with four independent external time servers.

To aid set up, the Pre-Cluster API of POST /cluster supports a list of NTP time servers using either the host name, IPv4 address, or IPv6 address.

You can enhance time security by acquiring private keys from external time servers, recording those keys and configuring the entries that match the external time servers to use those keys.

To use NTP symmetric authentication keys (keys), the shared private key must be recorded first using the /cluster/ntp/keys API associated with the server and enabled to be used.

APIs

There are three sets of APIs. The most basic set is part of the /api/cluster APIs, in which a set of NTP servers are provided.

The next two sets are used to manage the NTP servers in more detail and optionally record keys to enable NTP symmetric authentication.

[/api/cluster](#)

More details can be found under the documentation for [/api/cluster](#) . This API supports a list of NTP servers to start with. It does not take any individual configuration values for the NTP servers themselves.

[/api/cluster/ntp/servers](#)

You can use this API for a more detailed configuration of NTP servers. You must use this API to set and enable NTP symmetric authentication keys.

[/api/cluster/ntp/keys](#)

You can use this API to manage shared NTP symmetric keys that are provided by the remote NTP time server by using the key identifier (ID), type of key, and the private shared key.

Manage cluster NTP keys

Cluster NTP keys endpoint overview

Overview

You can configure NTP to use shared private keys between ONTAP and trusted external NTP time servers.

You acquire the keys from the external NTP time servers and individual entries created for each unique key. You can use the `/cluster/ntp/servers` API to associate a key with an external NTP time server used by ONTAP and enable authentication.

Fields used for adding an NTP shared key

The required fields are:

- `id`
- `digest_type`
- `secret_key`

Example

```
# Body
create_ntp_key.txt(body):
{
  "id": 10,
  "digest_type": "sha1",
  "value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/ntp/keys" -d
"@create_ntp_key.txt"
```

Retrieve NTP symmetric authentication keys

GET `/cluster/ntp/keys`

Introduced In: 9.7

Retrieves the collection of NTP symmetric authentication keys known by ONTAP that are uniquely indexed by an identifier.

Related ONTAP commands

- `cluster time-service ntp key show`

Learn more

- [DOC /cluster/ntp/keys](#)

Parameters

Name	Type	In	Required	Description
value	string	query	False	Filter by value
id	integer	query	False	Filter by id
digest_type	string	query	False	Filter by digest_type
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records.
records	array[ntp_key]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "1",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "digest_type": "sha1",
      "id": "10",
      "value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

ntp_key

Name	Type	Description
_links	_links	
digest_type	string	The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header.
id	integer	NTP symmetric authentication key identifier or index number (ID). This ID is included in the NTP cryptographic hash encoded header.
value	string	<p>A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.</p> <p>The current expected length is 40 characters.</p> <p>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data.</p>

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create an NTP symmetric authentication key entry

POST `/cluster/ntp/keys`

Introduced In: 9.7

Creates an NTP symmetric authentication key entry including the type of key using an unused identifier or index number (ID).

Required properties

- `id` - Shared symmetric key number (ID).
- `digest_type` - Shared private key cryptographic hash type.
- `value` - Value of shared private key.

Related ONTAP commands

- `cluster time-service ntp key create`

Learn more

- [DOC /cluster/ntp/keys](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
digest_type	string	The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header.
id	integer	NTP symmetric authentication key identifier or index number (ID). This ID is included in the NTP cryptographic hash encoded header.
value	string	<p>A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server. The current expected length is 40 characters.</p> <p>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data.</p>

Example request

```
{
  "digest_type": "sha1",
  "id": "10",
  "value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2097187	Invalid value for an NTP symmetric authentication key. A SHA1 key must be exactly 40 hexadecimal digits.
2097189	Too many NTP keys have been configured.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

ntp_key

Name	Type	Description
digest_type	string	The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header.
id	integer	NTP symmetric authentication key identifier or index number (ID). This ID is included in the NTP cryptographic hash encoded header.
value	string	<p>A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.</p> <p>The current expected length is 40 characters.</p> <p>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data.</p>

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete an NTP key

DELETE /cluster/ntp/keys/{id}

Introduced In: 9.7

Deletes an NTP key.

Related ONTAP commands

- `cluster time-service ntp key delete`

Learn more

- [DOC /cluster/ntp/keys](#)

Parameters

Name	Type	In	Required	Description
id	integer	path	True	Key identifier

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2097186	The key cannot be deleted because it is being used by an NTP server.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the NTP symmetric authentication key details

GET /cluster/ntp/keys/{id}

Introduced In: 9.7

Retrieves the details of a specific NTP symmetric authentication key by numeric identifier or index (ID).

Related ONTAP commands

- `cluster time-service ntp key show`

Learn more

- [DOC /cluster/ntp/keys](#)

Parameters

Name	Type	In	Required	Description
id	integer	path	True	Key identifier
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
digest_type	string	The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header.
id	integer	NTP symmetric authentication key identifier or index number (ID). This ID is included in the NTP cryptographic hash encoded header.

Name	Type	Description
value	string	<p>A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.</p> <p>The current expected length is 40 characters.</p> <p>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data.</p>

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "digest_type": "sha1",
  "id": "10",
  "value": "da39a3ee5e6b4b0d3255bfef95601890afd80709"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update NTP symmetric authentication key details

PATCH `/cluster/ntp/keys/{id}`

Introduced In: 9.7

Updates the details of a specific NTP symmetric authentication key by numeric identifier or index (ID).

Required properties

- `digest_type` - Shared private key cryptographic hash type.
- `value` - Value of shared private key.

Related ONTAP commands

- `cluster time-service ntp key modify`

Learn more

- [DOC /cluster/ntp/keys](#)

Parameters

Name	Type	In	Required	Description
id	integer	path	True	Key identifier

Request Body

Name	Type	Description
digest_type	string	The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header.
value	string	<p>A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.</p> <p>The current expected length is 40 characters.</p> <p>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data.</p>

Example request

```
{
  "digest_type": "sha1",
  "value": "da39a3ee5e6b4b0d3255bfeef95601890afd80709"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2097187	An invalid SHA1 key was provided.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

ntp_key

Name	Type	Description
digest_type	string	The type of cryptographic hash used to create and verify the NTP's message authentication code appended to each NTP packet header.
value	string	<p>A hexadecimal digit string that represents the cryptographic key that is shared with the remote NTP server.</p> <p>The current expected length is 40 characters.</p> <p>Use the cryptographic key and key ID to create a unique hash value used to authenticate the rest of the NTP data.</p>

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Manage cluster NTP servers

Cluster NTP servers endpoint overview

Overview

You can use this API to add external NTP servers to a cluster, update the configuration, use NTP keys, and retrieve the current NTP server configuration.

Adding an NTP server to a cluster

To add an NTP server to a cluster, issue a POST `/cluster/ntp/servers` request.

Fields used for adding an NTP server

Except for the name of the NTP server (host name or IP address), which is specified by the server, all fields are optional:

- `version`
- `key`

If the key is provided in POST, `authentication_enabled` is set to `true` by default.

Examples

Adding an NTP server

```
# Body
add_ntp_server.txt (body) :
{
  "server": "time.nist.gov"
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/ntp/servers" -d
"@add_ntp_server.txt"
```

Adding an NTP server with an authentication key

```
# Body
add_authenticated_ntp_server.txt (body) :
{
  "server": "time.nist.gov",
  "key": { "id": 10 }
}

# Request
curl -X POST "https://<mgmt-ip>/api/cluster/ntp/servers" -d
"@add_authenticated_ntp_server.txt"
```

Enabling a previously configured shared key (ID, type, and value) for an NTP server

A combination of key number or identifier (ID), type of key, and shared key value is created with `/api/cluster/ntp/keys`.

This operation will validate the NTP authentication works.

```
# Body
enable_shared_key.txt (body) :
{
  "key": { "id": 10 },
  "authentication_enabled": true
}

# Request
curl -X PATCH "https://<mgmt-ip>/api/cluster/ntp/servers/time.nist.gov" -d
"@enable_shared_key.txt"
```

Retrieve external NTP time servers

GET `/cluster/ntp/servers`

Introduced In: 9.7

Retrieves the collection of external NTP time servers ONTAP uses for time adjustment and correction.

Related ONTAP commands

- `cluster time-service ntp server show`

Learn more

- [DOC /cluster/ntp/servers](#)

Parameters

Name	Type	In	Required	Description
key.id	integer	query	False	Filter by key.id
version	string	query	False	Filter by version
authentication_enabled	boolean	query	False	Filter by authentication_enabled
server	string	query	False	Filter by server
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records.
records	array[ntp_server]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "3",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_enabled": 1,
      "key": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": "10"
      },
      "server": "time.nist.gov",
      "version": "auto"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

ntp_key_reference

Name	Type	Description
_links	_links	
id	integer	NTP symmetric authentication key identifier or index number (ID). This ID, the type of cryptographic hash, and the cryptographic hash value are all provided by the remote NTP server.

ntp_server

Name	Type	Description
_links	_links	
authentication_enabled	boolean	Set NTP symmetric authentication on (true) or off (false).
key	ntp_key_reference	
server	string	NTP server host name, IPv4, or IPv6 address.

Name	Type	Description
version	string	NTP protocol version for server. Valid versions are 3, 4, or auto.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Validate an external NTP time server

POST /cluster/ntp/servers

Introduced In: 9.7

Validates the provided external NTP time server for usage and configures ONTAP so that all nodes in the cluster use it.

The required fields are:

- `server`

Default property values

If not specified in POST, the following default property values are assigned:

- `version` - auto
- `key` - not set

If the key is provided in POST, `authentication_enabled` is set to `true` by default.

Related ONTAP commands

- `cluster time-service ntp server create`

Learn more

- [DOC /cluster/ntp/servers](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
key	ntp_key_reference	
server	string	NTP server host name, IPv4, or IPv6 address.
version	string	NTP protocol version for server. Valid versions are 3, 4, or auto.

Example request

```
{
  "authentication_enabled": 1,
  "key": {
    "id": "10"
  },
  "server": "time.nist.gov",
  "version": "auto"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2097163	NTP server IPv4 address was invalid.
2097164	NTP server IPv6 address was invalid.
2097165	Cannot resolve NTP server name.
2097166	NTP server address query returned no valid IP addresses.
2097167	Failed to connect to NTP server.
2097169	NTP server provided was not synchronized with a clock or another NTP server.
2097174	NTP server provided had too high of root distance.
2097177	NTP server provided an invalid stratum.
2097179	Too many NTP servers have been configured.
2097181	NTP server address was invalid. It is a special purpose address such as loopback, multicast, or broadcast address.
2097182	NTP server address was invalid. The address is neither an IPv4 or IPv6.
2097183	NTP symmetric key authentication cannot be used for a node not in a cluster.
2097185	NTP key authentication failed for the provided key.
2097193	An unknown NTP key was provided.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

ntp_key_reference

Name	Type	Description
id	integer	NTP symmetric authentication key identifier or index number (ID). This ID, the type of cryptographic hash, and the cryptographic hash value are all provided by the remote NTP server.

ntp_server

Name	Type	Description
key	ntp_key_reference	
server	string	NTP server host name, IPv4, or IPv6 address.
version	string	NTP protocol version for server. Valid versions are 3, 4, or auto.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete an external NTP server

```
DELETE /cluster/ntp/servers/{server}
```

Introduced In: 9.7

Deletes an external NTP server used by ONTAP.

Related ONTAP commands

- `cluster time-service ntp server delete`

Learn more

- [DOC /cluster/ntp/servers](#)

Parameters

Name	Type	In	Required	Description
server	string	path	True	Server address or host name

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve an external NTP server configuration

GET /cluster/ntp/servers/{server}

Introduced In: 9.7

Retrieves the configuration of an external NTP server used by ONTAP.

Related ONTAP commands

- `cluster time-service ntp server show`

Learn more

- [DOC /cluster/ntp/servers](#)

Parameters

Name	Type	In	Required	Description
server	string	path	True	NTP server host name, IPv4, or IPv6 address.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
authentication_enabled	boolean	Set NTP symmetric authentication on (true) or off (false).
key	ntp_key_reference	
server	string	NTP server host name, IPv4, or IPv6 address.
version	string	NTP protocol version for server. Valid versions are 3, 4, or auto.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_enabled": 1,
  "key": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "id": "10"
},
"server": "time.nist.gov",
"version": "auto"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

ntp_key_reference

Name	Type	Description
_links	_links	
id	integer	NTP symmetric authentication key identifier or index number (ID). This ID, the type of cryptographic hash, and the cryptographic hash value are all provided by the remote NTP server.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update an NTP server configuration after validation

PATCH /cluster/ntp/servers/{server}

Introduced In: 9.7

Updates the configuration of an NTP server used by the ONTAP cluster after validation.

Patchable fields are:

- version
- key.id
- authentication_enabled

If `authentication_enabled` is modified to `false`, the associated NTP key is removed from the server instance.

If `authentication_enabled` is modified to `true`, you must provide an NTP key ID in the PATCH body.

Related ONTAP commands

- `cluster time-service ntp server modify`

Learn more

- [DOC /cluster/ntp/servers](#)

Parameters

Name	Type	In	Required	Description
server	string	path	True	Server address or host name

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
authentication_enabled	boolean	Set NTP symmetric authentication on (true) or off (false).
key	ntp_key_reference	
version	string	NTP protocol version for server. Valid versions are 3, 4, or auto.

Example request

```
{
  "authentication_enabled": 1,
  "key": {
    "id": "10"
  },
  "version": "auto"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2097163	NTP server address was invalid.
2097164	NTP server address was invalid.
2097165	Could not resolve NTP server hostname.
2097166	NTP server address query returned no valid IP addresses.
2097167	Failed to connect to NTP server.

Error Code	Description
2097169	NTP server provided was not synchronized.
2097174	NTP server provided had too high of root distance.
2097177	NTP server provided had an invalid stratum.
2097181	NTP server address was invalid.
2097182	NTP server address was invalid.
2097183	NTP symmetric key authentication cannot be used for a node not in a cluster.
2097185	NTP key authentication failed for the provided key.
2097188	An invalid key identifier was provided. Identifiers must be in the range from 1 to 65535.
2097193	An unknown key was provided.
2097194	The field "authentication_enabled" cannot be false when the field NTP key is given.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

ntp_key_reference

Name	Type	Description
id	integer	NTP symmetric authentication key identifier or index number (ID). This ID, the type of cryptographic hash, and the cryptographic hash value are all provided by the remote NTP server.

ntp_server

Name	Type	Description
authentication_enabled	boolean	Set NTP symmetric authentication on (true) or off (false).
key	ntp_key_reference	
version	string	NTP protocol version for server. Valid versions are 3, 4, or auto.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster peers

Cluster peers endpoint overview

Overview

Cluster peering allows administrators of ONTAP systems to establish relationships between two or more independent clusters. When a relationship exists between two clusters, the clusters can exchange user data and configuration information, and coordinate operations. The `/cluster/peers` endpoint supports create, get, modify, and delete operations using GET, PATCH, POST and DELETE HTTP requests.

Create a cluster peer

You can set up a new cluster peer relationship by issuing a POST request to `/cluster/peers`. Parameters in the POST body define the settings of the peering relationship. A successful POST request that succeeds in creating a peer returns HTTP status code "201", along with the details of the created peer, such as peer UUID, name, and authentication information. A failed POST request returns an HTTP error code along with a message indicating the reason for the error. This can include malformed requests and invalid operations.

Sample request

```
curl -X POST 'https://<mgmt-ip>/api/cluster/peers/' -d
'{"authentication":{"expiry_time":"12/25/2018
12:34:56","generate_passphrase":true}}'
```

Examples

```

# Create - no params
body = {}

# Creating with a peer address and a passphrase
body =
{
  "remote":
  {
    "ip_addresses":["1.2.3.4"]
  }
}

# Creating with a peer name and a generated passphrase that is true
body =
{
  "name":"cp_xyz123",
  "authentication":
  {
    "generate_passphrase":true
  }
}

# Creating with a name, a peer address, and a passphrase
body =
{
  "name":"cp_xyz123",
  "remote":
  {
    "ip_addresses": ["1.2.3.4"]
  },
  "authentication":
  {
    "passphrase":"xyz12345"
  }
}

# Creating with a proposed encryption protocol
body =
{
  "encryption":
  {
    "proposed":"tls-psk"
  }
}

```

Create local intercluster LIFs

The local cluster must have an intercluster LIF on each node for the correct operation of cluster peering. If no local intercluster LIFs exist, you can optionally specify LIFs to be created for each node in the local cluster.

These local interfaces, if specified, are created on each node before proceeding with the creation of the cluster peering relationship. Cluster peering relationships are not established if there is an error preventing the LIFs from being created.

After local interfaces have been created, do not specify them for subsequent cluster peering relationships.

Local LIF creation fields

- `local_network.ip_addresses` - List of IP addresses to assign, one per node in the local cluster.
- `local_network.netmask` - IPv4 mask or subnet mask length.
- `local_network.broadcast_domain` - Broadcast domain that is in use within the IPspace.
- `local_network.gateway` - The IPv4 or IPv6 address of the default router.

Additional information on network routes

When creating LIFs, the network route discovery mechanism might take additional time (1-5 seconds) to become visible in the network outside of the cluster. This delay in publishing the routes might cause an initial cluster peer "create" request to fail. This error disappears with a retry of the same request.

Example

This example shows the POST body when creating four intercluster LIFs on a 4-node cluster before creating a cluster peer relationship.

`cluster_peer_4_node.txt`:

```
{
  "local_network":
  {
    "interfaces": [
      {"ip_address": "1.2.3.4"},
      {"ip_address": "1.2.3.5"},
      {"ip_address": "1.2.3.6"}
    ],
    "netmask": "255.255.0.0",
    "broadcast_domain": "Default",
    "gateway": "1.2.0.1"
  },
  "remote.ip_addresses": ["1.2.9.9"],
  "authentication.passphrase": "xyz12345"
}
curl -X POST "https://<mgmt-ip>/api/cluster/peers" -d
"@cluster_peer_4_node.txt"
```

Note that "<mgmt-ip>" is replaced by the IP address of the cluster management interface, and the body is read from the specified text file containing the fields for the new peering relationship and local interfaces.</mgmt-ip>

Retrieve a cluster peer

You can retrieve peers in a cluster by issuing a GET request to `/cluster/peers`. It is also possible to retrieve a specific peer when qualified by its UUID to `/cluster/peers/{uuid}`.

A GET request might have no query parameters or a valid cluster UUID. The former retrieves all records while the latter retrieves the record for the cluster peer with that UUID.

The following fields are used for retrieving a cluster peer.

Required fields

There are no required fields for GET requests.

Optional fields

The following fields are optional for GET requests

- `UUID` - UUID of the cluster peer.

Examples

```
curl -X GET "https://<mgmt-ip>/api/cluster/peers/"
```

```
curl -X GET "https://<mgmt-ip>/api/cluster/peers/{uuid}"
```

```
curl -X GET "https://<mgmt-ip>/api/cluster/peers/{uuid}?fields=*"
```

Update a cluster peer

You can update a cluster peer relationship by issuing a PATCH request to `/cluster/peers/{uuid}`. As in the CLI mode, you can toggle the proposed encryption protocol, update the passphrase, or specify a new set of stable addresses. All PATCH requests take the parameters that are to be updated in the request body. If the `generate_passphrase` is "true", the passphrase is returned in the PATCH response.

This following fields highlight the parameters that control the modification of an existing cluster peering relationship.

Required fields

A PATCH request with an empty body has no effect on the cluster peer instance. All other fields and the combinations in which they are valid are indicated below:

- `encryption_proposed` - Toggle the proposed encryption protocol (from "none" to "tls-psk" or otherwise). Authentication must be "true" and a passphrase must be present in body.
- `passphrase`
- `passphrase` or `generate passphrase`
- `remote.ip_addresses`

Optional fields

- `expiration time` - Set the expiration time of the passphrase.

Examples


```

# Updating with an empty body
body = {}

# Updating the proposed encryption protocol from tls-psk to none
body =
{
  "authentication":
    {
      "passphrase": "xyz12345",
      "in_use": "ok"
    },
  "encryption":
    {
      "proposed": "none"
    }
}

# Updating the passphrase
body =
{
  "authentication":
    {
      "passphrase": "xyz12345",
      "in_use": "ok"
    }
}

# Setting an auto-generated passphrase
body =
{
  "authentication":
    {
      "generate_passphrase": true,
      "in_use": "ok"
    }
}

# Updating remote IP addresses
body =
{
  "remote":
    {
      "ip_addresses": ["10.224.65.30"]
    }
}

```

Sample requests

```
# Setting a passphrase
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/73123071-d0b9-11e8-a686-005056a7179a' -d
'{"authentication":{"passphrase":"xyz12345","in_use":"ok"}}'

# Updating a peer address
curl -X PATCH 'https://<mgmt-ip>/api/cluster/peers/73123071-d0b9-11e8-a686-005056a7179a' -d '{"remote":{"ip_addresses":["1.2.3.4"]}}'
```

Delete a cluster peer

You can delete a cluster peer using the HTTP DELETE request.

Required fields

Perform all delete operations on a valid peer UUID. Deleting an invalid peer returns "HTTP 404", which indicates an error.

Optional fields

The DELETE operation has no optional fields.

Request format

DELETE "https://<mgmt-ip>/api/cluster/peers/{uuid}"</mgmt-ip>

Example

The following request deletes a peer with peer UUID "8becc0d4-c12c-11e8-9ceb-005056bbd143".

```
curl -X DELETE "https://<mgmt-ip>/api/cluster/peers/8becc0d4-c12c-11e8-9ceb-005056bbd143"
```

Retrieve cluster peers

GET /cluster/peers

Introduced In: 9.6

Retrieves the collection of cluster peers.

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cluster_peer]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "interfaces": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication": {
        "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
        "in_use": "string",
        "passphrase": "string",
        "state": "string"
      },
      "encryption": {
        "proposed": "string",
        "state": "string"
      },
      "initial_allowed_svms": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      ],
      "ipospace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    }
  ]
}
```

```

    }
  },
  "name": "exchange",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"name": "cluster2",
"peer_applications": [
  "snapmirror",
  "flexcache"
],
"remote": {
  "ip_addresses": [
    "10.10.10.7"
  ],
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25T11:20:13Z"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"version": {
  "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
  "generation": "9",
  "major": "4",
  "minor": "0"
}
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
interfaces	href	
self	href	

authentication

Name	Type	Description
expiry_time	string	The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour.
generate_passphrase	boolean	Auto generate a passphrase when true.
in_use	string	
passphrase	string	A password to authenticate the cluster peer relationship.
state	string	

encryption

Name	Type	Description
proposed	string	
state	string	

_links

Name	Type	Description
self	href	

initial_allowed_svms

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipspace

The IPspace of the local intercluster LIFs.

Name	Type	Description
_links	_links	
name	string	IPspace name
uuid	string	IPspace UUID

interfaces

Name	Type	Description
ip_address	string	IPv4 or IPv6 address

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
broadcast_domain	string	Broadcast domain that is in use within the IPspace.
gateway	string	The IPv4 or IPv6 address of the default router.
interfaces	array[interfaces]	
netmask	string	IPv4 mask or netmask length.

remote

Name	Type	Description
ip_addresses	array[string]	The IPv4 addresses, IPv6 addresses, or hostnames of the peers.
name	string	The name of the remote cluster.
serial_number	string	The serial number of the remote cluster.

status

Name	Type	Description
state	string	
update_time	string	The last time the state was updated.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

cluster_peer

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.

Name	Type	Description
ipspace	ipspace	The IPspace of the local intercluster LIFs.
name	string	Optional name for the cluster peer relationship. By default, it is the name of the remote cluster.
peer_applications	array[string]	Peering applications against which allowed SVMs are configured.
remote	remote	
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Create a peering relationship

POST /cluster/peers

Introduced In: 9.6

Creates a peering relationship and, optionally, the IP interfaces it will use. There are two methods used to create a peering relationship:

- Provide a remote IP address - Used when creating a new cluster peer relationship with a specific remote cluster. This requires at least one remote intercluster IP address from the remote cluster.
- Do not provide a remote IP address - Used when the remote IP address is not provided and when the storage system is ready to accept peering requests from foreign clusters.

Required properties

- `remote_ip_addresses` - Addresses of the remote peers. The local peer must be able to reach and connect to these addresses for the request to succeed in creating a peer. Only required when creating a peering relationship by providing a remote IP address.
- Either set `generate_passphrase` to "true" or provide a passphrase in the body of the request. Only one of these options is required.

Recommended optional properties

- `name` - Name of the peering relationship or name of the remote peer.
- `passphrase` - User generated passphrase for use in authentication.
- `generate_passphrase` (true/false) - When "true", ONTAP automatically generates a passphrase to authenticate cluster peers.
- `ipspace` - IPspace of the local intercluster LIFs. Assumes Default IPspace if not provided.
- `initial_allowed_svms` - Local SVMs allowed to peer with the peer cluster's SVMs. Can be modified until the remote cluster accepts this cluster peering relationship.
- `local_network` - Fields to create a local intercluster LIF.
- `expiry_time` - Duration in ISO 8601 format for which the user-supplied or auto-generated passphrase is valid. Expiration time must not be greater than seven days into the future. ISO 8601 duration format is "PnDTnHnMnS" or "PnW" where n is a positive integer. The "nD", "nH", "nM" and "nS" fields can be dropped if zero. "P" must always be present and "T" must be present if there are any hours, minutes, or seconds fields.
- `encryption_proposed` (none/tls-psk) - Encryption mechanism of the communication channel between the two peers.
- `peer_applications` - SVM peering applications (SnapMirror, FlexCache or both) for which the SVM peering relationship is set up.

Additional information

As with creating a cluster peer through the CLI, the combinations of options must be valid in order for the create operation to succeed. The following list shows the combinations that will succeed and those that will fail:

- A passphrase only (fail)
- A peer IP address (fail)
- A passphrase with an expiration time > 7 days into the future (fail)
- A peer IP address and a passphrase (OK)
- generate_passphrase=true (OK)
- Any proposed encryption protocol (OK)
- An IPspace name or UUID (OK)
- A passphrase, peer IP address, and any proposed encryption protocol (OK)
- A non empty list of initial allowed SVM peer names or UUIDs. (OK)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs.
local_network	local_network	Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
name	string	Optional name for the cluster peer relationship. By default, it is the name of the remote cluster.
peer_applications	array[string]	Peering applications against which allowed SVMs are configured.
remote	remote	
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

Example request

```
{
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "string",
    "passphrase": "string",
    "state": "string"
  },
  "encryption": {
    "proposed": "string",
    "state": "string"
  },
  "initial_allowed_svms": [
    {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "ipspace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local_network": {
    "broadcast_domain": "bd1",
    "gateway": "10.1.1.1",
    "interfaces": [
      {
        "ip_address": "10.10.10.7"
      }
    ],
    "netmask": "255.255.0.0"
  },
  "name": "cluster2",
  "peer_applications": [
    "snapmirror",
    "flexcache"
  ],
  "remote": {
    "ip_addresses": [
      "10.10.10.7"
    ],
    "name": "cluster2",
    "serial_number": "4048820-60-9"
  },
  "status": {
```

```
"state": "available",
"update_time": "2017-01-25T11:20:13Z"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 201, Created

Name	Type	Description
authentication	authentication	
ip_address	string	IPv4 or IPv6 address
name	string	Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers.

Example response

```
{
  "authentication": {
    "expiry_time": "2017-01-25T11:20:13Z",
    "passphrase": "string"
  },
  "ip_address": "10.10.10.7",
  "name": "cluster2"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1966366	The system SVM of the cluster IPspace hosts cluster LIFs only.
4653365	IPspaces are unavailable with cluster peering: {ipspace}.
4656069	Specifying a passphrase without remote IP addresses is not supported.
4656070	The encryption protocol is meaningful only with authenticated cluster peer relationships.
4656071	Cannot peer with a cluster bearing the same name as the local cluster.
4656072	The name must conform to the same rules as a cluster name.
4656074	Cannot check whether all nodes of this cluster support encryption.
4656075	Cannot specify encryption: this operation requires an ECV of 9.6.0 or later.
4656077	Specify either remote IP addresses or generate_passphrase.
4656079	No cluster nodes were found. Check your cluster configuration.
4656081	Creating an intercluster LIF requires a list of local IP addresses.
4656085	Cannot create an intercluster LIF with an empty list of local IP addresses.
4656086	Creating an intercluster LIF requires a broadcast domain that is in use within the IPspace.
4656087	The number of local intercluster IP addresses must be less than or equal to the number of available nodes.
4656088	Found no ports matching the IPspace and the broadcast domain.
4656089	Found no matching entry for IPspace.
4656090	The given IPspace differs from the IPspace entry found.
4656091	Creating an intercluster LIF requires a subnet mask or a subnet mask length.
4656096	Creating an intercluster LIF requires an IPv4 or IPv6 address of the default router.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

authentication

Name	Type	Description
expiry_time	string	The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour.
generate_passphrase	boolean	Auto generate a passphrase when true.
in_use	string	
passphrase	string	A password to authenticate the cluster peer relationship.
state	string	

encryption

Name	Type	Description
proposed	string	
state	string	

initial_allowed_svms

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipspace

The IPspace of the local intercluster LIFs.

Name	Type	Description
name	string	IPspace name

Name	Type	Description
uuid	string	IPspace UUID

interfaces

Name	Type	Description
ip_address	string	IPv4 or IPv6 address

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
broadcast_domain	string	Broadcast domain that is in use within the IPspace.
gateway	string	The IPv4 or IPv6 address of the default router.
interfaces	array[interfaces]	
netmask	string	IPv4 mask or netmask length.

remote

Name	Type	Description
ip_addresses	array[string]	The IPv4 addresses, IPv6 addresses, or hostnames of the peers.
name	string	The name of the remote cluster.
serial_number	string	The serial number of the remote cluster.

status

Name	Type	Description
state	string	
update_time	string	The last time the state was updated.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

cluster_peer

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs.
local_network	local_network	Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.
name	string	Optional name for the cluster peer relationship. By default, it is the name of the remote cluster.
peer_applications	array[string]	Peering applications against which allowed SVMs are configured.
remote	remote	

Name	Type	Description
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

authentication

Name	Type	Description
expiry_time	string	The date and time the passphrase will expire. The default expiry time is one hour.
passphrase	string	A password to authenticate the cluster peer relationship.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a cluster peer

DELETE /cluster/peers/{uuid}

Introduced In: 9.6

Deletes a cluster peer.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cluster peer relationship UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
4663070	Unable to delete cluster peer relationship due to an ongoing Vserver migration.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a cluster peer instance

GET /cluster/peers/{uuid}

Introduced In: 9.6

Retrieves a specific cluster peer instance.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cluster peer relationship UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs.
name	string	Optional name for the cluster peer relationship. By default, it is the name of the remote cluster.
peer_applications	array[string]	Peering applications against which allowed SVMs are configured.
remote	remote	
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.
version	version	This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Example response

```
{
  "_links": {
    "interfaces": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "string",
    "passphrase": "string",
    "state": "string"
  },
  "encryption": {
    "proposed": "string",
    "state": "string"
  },
  "initial_allowed_svms": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "cluster2",
  "peer_applications": [
    "snapmirror",
    "flexcache"
  ],
}
```



```

"remote": {
  "ip_addresses": [
    "10.10.10.7"
  ],
  "name": "cluster2",
  "serial_number": "4048820-60-9"
},
"status": {
  "state": "available",
  "update_time": "2017-01-25T11:20:13Z"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"version": {
  "full": "NetApp Release 9.4.0: Sun Nov 05 18:20:57 UTC 2017",
  "generation": "9",
  "major": "4",
  "minor": "0"
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
interfaces	href	
self	href	

authentication

Name	Type	Description
expiry_time	string	The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour.
generate_passphrase	boolean	Auto generate a passphrase when true.
in_use	string	
passphrase	string	A password to authenticate the cluster peer relationship.
state	string	

encryption

Name	Type	Description
proposed	string	
state	string	

_links

Name	Type	Description
self	href	

initial_allowed_svms

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipspace

The IPspace of the local intercluster LIFs.

Name	Type	Description
_links	_links	
name	string	IPspace name
uuid	string	IPspace UUID

interfaces

Name	Type	Description
ip_address	string	IPv4 or IPv6 address

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
broadcast_domain	string	Broadcast domain that is in use within the IPspace.
gateway	string	The IPv4 or IPv6 address of the default router.
interfaces	array[interfaces]	
netmask	string	IPv4 mask or netmask length.

remote

Name	Type	Description
ip_addresses	array[string]	The IPv4 addresses, IPv6 addresses, or hostnames of the peers.

Name	Type	Description
name	string	The name of the remote cluster.
serial_number	string	The serial number of the remote cluster.

status

Name	Type	Description
state	string	
update_time	string	The last time the state was updated.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Update a cluster peer instance

PATCH /cluster/peers/{uuid}

Introduced In: 9.6

Updates a cluster peer instance.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Cluster peer relationship UUID

Request Body

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs.
name	string	Optional name for the cluster peer relationship. By default, it is the name of the remote cluster.
peer_applications	array[string]	Peering applications against which allowed SVMs are configured.
remote	remote	
status	status	

Name	Type	Description
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

Example request

```
{
  "authentication": {
    "expiry_time": "P1DT2H3M4S or '2017-01-25T11:20:13Z'",
    "in_use": "string",
    "passphrase": "string",
    "state": "string"
  },
  "encryption": {
    "proposed": "string",
    "state": "string"
  },
  "initial_allowed_svms": [
    {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ],
  "ipspace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "cluster2",
  "peer_applications": [
    "snapmirror",
    "flexcache"
  ],
  "remote": {
    "ip_addresses": [
      "10.10.10.7"
    ],
    "name": "cluster2",
    "serial_number": "4048820-60-9"
  },
  "status": {
    "state": "available",
    "update_time": "2017-01-25T11:20:13Z"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```


Response

Status: 200, Ok

Name	Type	Description
authentication	authentication	
ip_address	string	IPv4 or IPv6 address
name	string	Optional name for the cluster peer relationship. By default, it is the name of the remote cluster, or a temporary name might be autogenerated for anonymous cluster peer offers.

Example response

```
{
  "authentication": {
    "expiry_time": "2017-01-25T11:20:13Z",
    "passphrase": "string"
  },
  "ip_address": "10.10.10.7",
  "name": "cluster2"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
4653261	Error finding IPspace.
4655058	Expiration time cannot be more than 7 days in the future.
4656070	The encryption protocol is meaningful only with authenticated cluster peer relationships.
4656072	The name must conform to the same rules as a cluster name.

Error Code	Description
4656073	Changing the encryption state requires the refreshing of the authentication passphrase.
4656075	Cannot specify encryption: this operation requires an ECV of ONTAP 9.6.0 or later.
4656076	Cluster peer modify was attempted with mismatched IPv4 and IPv6 addresses.
4656081	The remote IP address list is empty.
4656082	Specify either a passphrase or "-generate-passphrase".
4656083	Cannot auto-generate a passphrase when "generate-passphrase" is false. Modifying a passphrase using an auto-generated passphrase requires "generate-passphrase" be true.
4656084	Passphrase can only be modified with an authenticated cluster peer relationship.
4656092	Cluster peer modify was attempted with a host name that did not resolve to an IPv4 or IPv6 address.
4656095	The address family of the specified peer addresses is not valid in this IPspace. Use /api/network/interfaces/ to verify that required LIFs are present and operational on each cluster node.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

authentication

Name	Type	Description
expiry_time	string	The time when the passphrase will expire, in ISO 8601 duration format or date and time format. The default is 1 hour.
generate_passphrase	boolean	Auto generate a passphrase when true.
in_use	string	
passphrase	string	A password to authenticate the cluster peer relationship.
state	string	

encryption

Name	Type	Description
proposed	string	
state	string	

initial_allowed_svms

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipspace

The IPspace of the local intercluster LIFs.

Name	Type	Description
name	string	IPspace name

Name	Type	Description
uuid	string	IPspace UUID

interfaces

Name	Type	Description
ip_address	string	IPv4 or IPv6 address

local_network

Cluster peering requires an intercluster LIF on each local node. These can be optionally created by specifying a list of IP addresses corresponding to each node.

Name	Type	Description
broadcast_domain	string	Broadcast domain that is in use within the IPspace.
gateway	string	The IPv4 or IPv6 address of the default router.
interfaces	array[interfaces]	
netmask	string	IPv4 mask or netmask length.

remote

Name	Type	Description
ip_addresses	array[string]	The IPv4 addresses, IPv6 addresses, or hostnames of the peers.
name	string	The name of the remote cluster.
serial_number	string	The serial number of the remote cluster.

status

Name	Type	Description
state	string	
update_time	string	The last time the state was updated.

version

This returns the cluster version information. When the cluster has more than one node, the cluster version is equivalent to the lowest of generation, major, and minor versions on all nodes.

Name	Type	Description
full	string	The full cluster version string.
generation	integer	The generation portion of the version.
major	integer	The major portion of the version.
minor	integer	The minor portion of the version.

cluster_peer

Name	Type	Description
_links	_links	
authentication	authentication	
encryption	encryption	
initial_allowed_svms	array[initial_allowed_svms]	The local SVMs allowed to peer with the peer cluster's SVMs. This list can be modified until the remote cluster accepts this cluster peering relationship.
ipspace	ipspace	The IPspace of the local intercluster LIFs.
name	string	Optional name for the cluster peer relationship. By default, it is the name of the remote cluster.
peer_applications	array[string]	Peering applications against which allowed SVMs are configured.
remote	remote	
status	status	
uuid	string	UUID of the cluster peer relationship. For anonymous cluster peer offers, the UUID will change when the remote cluster accepts the relationship.

authentication

Name	Type	Description
expiry_time	string	The date and time the passphrase will expire. The default expiry time is one hour.
passphrase	string	A password to authenticate the cluster peer relationship.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster schedules

Cluster schedules endpoint overview

Overview

You can use the `/cluster/schedules` API to view, create, and modify job schedules in a cluster.

Retrieving a job schedule

You can retrieve job schedules by issuing a GET request to `/cluster/schedules`. It is also possible to retrieve a specific schedule when qualified by its UUID to `/cluster/schedules/{uuid}`. You can apply queries on fields to retrieve all schedules that match the combined query.

Example

```
# The API:
/api/cluster/schedules/

# The call:
curl -X GET 'https://<mgmt-ip>/api/cluster/schedules?type=interval'

# The response:
{
  "records": [
    {
      "uuid": "08ceae53-0158-11e9-a82c-005056bb4301",
      "name": "RepositoryBalanceMonitorJobSchedule",
      "type": "interval",
      "interval": "PT10M",
      "_links": {
        "self": {
          "href": "/api/cluster/schedules/08ceae53-0158-11e9-a82c-005056bb4301"
        }
      }
    },
    {
      "uuid": "0941e980-0158-11e9-a82c-005056bb4301",
      "name": "Balanced Placement Model Cache Update",
      "type": "interval",
      "interval": "PT7M30S",
      "_links": {
        "self": {
          "href": "/api/cluster/schedules/0941e980-0158-11e9-a82c-005056bb4301"
        }
      }
    },
    {
      "uuid": "0944b975-0158-11e9-a82c-005056bb4301",
      "name": "Auto Balance Aggregate Scheduler",
      "type": "interval",
      "interval": "PT1H",
      "_links": {
        "self": {
          "href": "/api/cluster/schedules/0944b975-0158-11e9-a82c-005056bb4301"
        }
      }
    }
  ]
}
```

```

    },
    {
      "uuid": "0c65f1fb-0158-11e9-a82c-005056bb4301",
      "name": "Application Templates ASUP Dump",
      "type": "interval",
      "interval": "P1D",
      "_links": {
        "self": {
          "href": "/api/cluster/schedules/0c65f1fb-0158-11e9-a82c-005056bb4301"
        }
      }
    }
  ],
  "num_records": 4,
  "_links": {
    "self": {
      "href": "/api/cluster/schedules?type=interval"
    }
  }
}

```



```
# The API:
/api/cluster/schedules/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/cluster/schedules/25312bd8-0158-11e9-a82c-005056bb4301'

# The response:
{
  "uuid": "25312bd8-0158-11e9-a82c-005056bb4301",
  "name": "monthly",
  "cluster": {
    "name": "rodan-tsundere",
    "uuid": "f3f9bbfa-0157-11e9-a82c-005056bb4301"
  },
  "type": "cron",
  "cron": {
    "minutes": [
      20
    ],
    "hours": [
      0
    ],
    "days": [
      1
    ]
  },
  "_links": {
    "self": {
      "href": "/api/cluster/schedules/25312bd8-0158-11e9-a82c-005056bb4301"
    }
  }
}
```

Creating a job schedule

You can create a job schedule by issuing a POST request to `/cluster/schedules` to a node in the cluster. For a successful request, the POST request returns a status code of 201.

Job schedules can be of either type "cron" or type "interval". A cron schedule is run at specific minutes within the hour, or hours of the day, days of the week, days of the month, or months of the year. An interval schedule runs repeatedly at fixed intervals.

Required fields

- name - Name of the job schedule

You are required to provide a "minutes" field for a cron schedule. An "interval" field is required for an interval schedule. Do not provide both a "cron" field and an "interval" field.
The schedule UUID is created by the system.

Cron schedule fields

- cron.minutes - Minutes within the hour (0 through 59)
- cron.hours - Hours of the day (0 through 23)
- cron.weekdays - Weekdays (0 through 6, where 0 is Sunday and 6 is Saturday.)
- cron.days - Days of the month (1 through 31)
- cron.months - Months of the year (1 through 12)

Interval schedule field

- interval - Length of time in ISO 8601 duration format.

Examples

Create an interval schedule with a 1-week interval

```
# The API:
/api/cluster/schedules
one_week_interval.txt:
{
  "name": "test_interval_1",
  "interval": "P1W"
}

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster/schedules" -d
"@one_week_interval.txt"

# The response of a successful POST is empty.
```

Create a cron schedule that runs daily at 12:05

```
# The API:
/api/cluster/schedules
daily_noon_job.txt:
{
  "name": "test_cron_1",
  "cron":
  {
    "minutes": [ 5 ],
    "hours": [ 12 ]
  }
}

# The call:
curl -X POST "https://<mgmt-ip>/api/cluster/schedules" -d
"@daily_noon_job.txt"

# The response of a successful POST is empty.
```

Optional fields

By default, the schedule is owned by the local cluster. In a MetroCluster configuration, you can specify the partner cluster if the local cluster is in the switchover state.

- `cluster.name` - Name of the cluster owning the schedule.
- `cluster.uuid` - UUID of the cluster owning the schedule.

Records field

You can create multiple schedules in one request by providing an array of named records with schedule entries. Each entry must follow the required and optional fields listed above.

Updating a job schedule

The following fields of an existing schedule can be modified:

- `cron.minutes`
- `cron.hours`
- `cron.weekdays`
- `cron.days`
- `cron.months`
- `interval`

Note that you cannot modify the name, cluster, and type of schedule. Also, you cannot modify a cron field of an interval schedule, or the interval field of a cron schedule. You can apply queries on fields to modify all schedules that match the combined query.

Examples

Modify an interval schedule with a 2-day and 5-minute interval

```
# The API:
/api/cluster/schedules/{uuid}
every_two_days_five_minutes.txt:
{
  "interval": "P2DT5M"
}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/cluster/schedules/{uuid}" -d
"@every_two_days_five_minutes.txt"

# The response of a successful PATCH is empty.
```

Modify a cron schedule to run Mondays at 2

```
# The API:
/api/cluster/schedules/{uuid}
monday_at_two.txt:
{
  "cron":
  {
    "hours": [ 2 ],
    "weekdays": [ 1 ]
  }
}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/cluster/schedules/{uuid}" -d
"@monday_at_two.txt"

# The response of a successful PATCH is empty.
```

Deleting a job schedule

You can delete job schedules based on their UUID. You can apply queries on fields to delete all schedules that match the combined query.

Example

```
# The API:
/api/cluster/schedules/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/schedules/{uuid}"

# The response of a successful DELETE of one schedule is empty.
```

```
# The API:
/api/cluster/schedules/

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/schedules/?name=test*"

# The response of a successful DELETE indicates the number of schedules
affected:
{
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/cluster/schedules?name=test*"
    }
  }
}
```

MetroCluster configurations

In a MetroCluster configuration, user-created schedules owned by the local cluster are replicated to the partner cluster. Likewise, user-created schedules owned by the partner cluster are replicated to the local cluster. The owning cluster for a particular schedule is shown in the "cluster.name" and "cluster.uuid" fields. Normally, only schedules owned by the local cluster can be created, modified, and deleted on the local cluster. However, when a MetroCluster configuration is in switchover, the cluster in switchover state can create, modify, and delete schedules owned by the partner cluster.

Retrieve schedules

GET /cluster/schedules

Introduced In: 9.6

Retrieves a schedule.

Parameters

Name	Type	In	Required	Description
name	string	query	False	Filter by name
interval	string	query	False	Filter by interval
uuid	string	query	False	Filter by uuid
cluster.uuid	string	query	False	Filter by cluster.uuid
cluster.name	string	query	False	Filter by cluster.name
cron.months	integer	query	False	Filter by cron.months
cron.days	integer	query	False	Filter by cron.days
cron.minutes	integer	query	False	Filter by cron.minutes
cron.weekdays	integer	query	False	Filter by cron.weekdays
cron.hours	integer	query	False	Filter by cron.hours
type	string	query	False	Filter by type
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[schedule]	

Example response


```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "1",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "cron": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      },
      "interval": "P1DT2H3M4S",
      "name": "string",
      "type": "string",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
459760	The schedule specified is not a valid schedule.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

cluster

The cluster that owns the schedule. Defaults to the local cluster.

Name	Type	Description
name	string	Cluster name
uuid	string	Cluster UUID

cron

Details for schedules of type cron.

Name	Type	Description
days	array[integer]	The days of the month the schedule runs. Leave empty for all.
hours	array[integer]	The hours of the day the schedule runs. Leave empty for all.
minutes	array[integer]	The minutes the schedule runs. Required on POST for a cron schedule.

Name	Type	Description
months	array[integer]	The months of the year the schedule runs. Leave empty for all.
weekdays	array[integer]	The weekdays the schedule runs. Leave empty for all.

schedule

Complete schedule information

Name	Type	Description
_links	_links	
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule name. Required in the URL or POST body.
type	string	Schedule type
uuid	string	Job schedule UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a schedule

POST /cluster/schedules

Introduced In: 9.6

Creates a schedule.

Required Fields

- name - Name of the job schedule.
You must provide a minutes field for a cron schedule and an interval field for an interval schedule. Do not provide both a cron field and an interval field.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule name. Required in the URL or POST body.

Name	Type	Description
type	string	Schedule type
uuid	string	Job schedule UUID

Example request

```
{
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": [
      "integer"
    ],
    "hours": [
      "integer"
    ],
    "minutes": [
      "integer"
    ],
    "months": [
      "integer"
    ],
    "weekdays": [
      "integer"
    ]
  },
  "interval": "P1DT2H3M4S",
  "name": "string",
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
458788	The schedule specified is not a valid schedule.
459760	The schedule specified is not a valid schedule.
459763	Schedule cannot be created locally using the remote cluster name as the owner.
459764	Cannot create a schedule with the same name as an existing schedule from the MetroCluster partner cluster but of a different schedule type.
460783	As this is a MetroCluster configuration and the local cluster is waiting for switchback, changes to non-system schedules are not allowed.
460784	An error occurred creating the remote cluster version of this schedule.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster

The cluster that owns the schedule. Defaults to the local cluster.

Name	Type	Description
name	string	Cluster name
uuid	string	Cluster UUID

cron

Details for schedules of type cron.

Name	Type	Description
days	array[integer]	The days of the month the schedule runs. Leave empty for all.
hours	array[integer]	The hours of the day the schedule runs. Leave empty for all.
minutes	array[integer]	The minutes the schedule runs. Required on POST for a cron schedule.
months	array[integer]	The months of the year the schedule runs. Leave empty for all.
weekdays	array[integer]	The weekdays the schedule runs. Leave empty for all.

schedule

Complete schedule information

Name	Type	Description
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule name. Required in the URL or POST body.
type	string	Schedule type
uuid	string	Job schedule UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a schedule

DELETE /cluster/schedules/{uuid}

Introduced In: 9.6

Deletes a schedule.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Schedule UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
459758	Cannot delete a job schedule that is in use. Remove all references to the schedule, and then try to delete again.
459761	Schedule cannot be deleted on this cluster because it is replicated from the remote cluster.
459762	The schedule cannot be deleted because it is a system-level schedule.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a schedule

GET /cluster/schedules/{uuid}

Introduced In: 9.6

Retrieves a schedule.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Schedule UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
cluster	cluster	The cluster that owns the schedule. Defaults to the local cluster.
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
name	string	Schedule name. Required in the URL or POST body.
type	string	Schedule type
uuid	string	Job schedule UUID

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "cron": {
    "days": [
      "integer"
    ],
    "hours": [
      "integer"
    ],
    "minutes": [
      "integer"
    ],
    "months": [
      "integer"
    ],
    "weekdays": [
      "integer"
    ]
  },
  "interval": "P1DT2H3M4S",
  "name": "string",
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

The cluster that owns the schedule. Defaults to the local cluster.

Name	Type	Description
name	string	Cluster name
uuid	string	Cluster UUID

cron

Details for schedules of type cron.

Name	Type	Description
days	array[integer]	The days of the month the schedule runs. Leave empty for all.
hours	array[integer]	The hours of the day the schedule runs. Leave empty for all.
minutes	array[integer]	The minutes the schedule runs. Required on POST for a cron schedule.
months	array[integer]	The months of the year the schedule runs. Leave empty for all.
weekdays	array[integer]	The weekdays the schedule runs. Leave empty for all.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a schedule

PATCH `/cluster/schedules/{uuid}`

Introduced In: 9.6

Updates a schedule. Note that you cannot modify a cron field of an interval schedule, or the interval field of a cron schedule.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Schedule UUID

Request Body

Name	Type	Description
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
type	string	Schedule type
uuid	string	Job schedule UUID

Example request

```
{
  "cron": {
    "days": [
      "integer"
    ],
    "hours": [
      "integer"
    ],
    "minutes": [
      "integer"
    ],
    "months": [
      "integer"
    ],
    "weekdays": [
      "integer"
    ]
  },
  "interval": "P1DT2H3M4S",
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
458788	The schedule specified is not a valid schedule.
459760	The schedule specified is not a valid schedule.
459761	Schedule cannot be modified on this cluster because it is replicated from the remote cluster.

Error Code	Description
460783	As this is a MetroCluster configuration and the local cluster is waiting for switchback, changes to non-system schedules are not allowed.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster

The cluster that owns the schedule. Defaults to the local cluster.

Name	Type	Description
name	string	Cluster name
uuid	string	Cluster UUID

cron

Details for schedules of type cron.

Name	Type	Description
days	array[integer]	The days of the month the schedule runs. Leave empty for all.
hours	array[integer]	The hours of the day the schedule runs. Leave empty for all.
minutes	array[integer]	The minutes the schedule runs. Required on POST for a cron schedule.
months	array[integer]	The months of the year the schedule runs. Leave empty for all.
weekdays	array[integer]	The weekdays the schedule runs. Leave empty for all.

schedule

Complete schedule information

Name	Type	Description
cron	cron	Details for schedules of type cron.
interval	string	An ISO-8601 duration formatted string.
type	string	Schedule type
uuid	string	Job schedule UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster software

Cluster software endpoint overview

Overview

You can use the ONTAP cluster software API to retrieve and display relevant information about a software profile, software packages collection, and software history collection. This API retrieves the information about all software packages present in the cluster, or a specific software package.

You can use the POST request to download a software package from an HTTP or FTP server. The PATCH request provides the option to upgrade the cluster software version. Select the `validate_only` field to validate the package before triggering the update. Set the `version` field to trigger the installation of the package in the cluster. You can pause, resume, or cancel any ongoing software upgrade by selecting `action`. You can use the DELETE request to remove a specific software package present in the cluster.

Retrieving software profile information

The following example shows how to retrieve software profile information. You can check the validation results after selecting the `validate_only` field. Upgrade progress information is available after an upgrade has started.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software?return_timeout=15" -H
"accept: application/hal+json"

# The response:
{
  "validation_results": [
    {
      "update_check": "NFS mounts",
      "status": "warning",
      "issue": {
        "message": "Use NFS hard mounts, if possible.",
      }
    },
    {
      "action": {
        "message": "Use NFS hard mounts, if possible.",
      }
    }
  ],
  "version": "9.5.0",
  "pending_version": "9.6.0",
  "nodes": [
    {
      "node": "sti70-vsim-ucs165n",
      "version": "9.5.0"
    }
  ],
  "metrocluster": {
    "progress_summary": {
      "message": "Update paused by user"
    },
    "progress_details": {
      "message": "Installing software image on cluster \"sti70-vsim-ucs165n_siteA\"."
    },
    "clusters": [
```

```

    {
      "name": "sti70-vsim-ucs165n_siteA",
      "uuid": "720f046c-4b13-11e9-9c34-005056ac5626",
      "estimated_duration": 3480,
      "elapsed_duration": 0,
      "state": "waiting"
    },
  ],
  ],
  "state": "in_progress",
  "start_time": "2018-05-21T09:53:04+05:30",
  "end_time": "2018-05-21T11:53:04+05:30",
  "estimated_time": 5220,
  "elapsed_time": 2140,
  "update_details": [
    {
      "phase": "Data ONTAP updates",
      "state": "in_progress",
      "estimated_duration": 4620,
      "elapsed_duration": 29,
      "node": {
        "name": "sti70-vsim-ucs165n"
      }
    }
  ],
  "status_details": [
    {
      "name": "do-download-job",
      "state": "completed",
      "issue": {
        "message": "Image update complete",
        "code": 0
      },
      "start_time": "2018-05-21T09:53:04+05:30",
      "end_time": "2018-05-21T11:53:04+05:30",
      "node": {
        "name": "sti70-vsim-ucs165n"
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster/software/"
    }
  }
}

```

Upgrading the software version

The following example shows how to upgrade cluster software. Set the `version` field to trigger the installation of the package. You can select the `validate_only` field to validate the package before the installation starts. Setting `skip_warning` as `true` ignores the validation warning before the installation starts. Setting the `action` field performs a `pause`, `resume`, or `cancel` operation on an ongoing upgrade. An upgrade can only be resumed if it is in the `paused` state. Setting `stabilize_minutes` allows each node a specified amount of time to stabilize after a reboot; the default is 8 minutes.

You can start the upgrade process at the cluster level. There are no options available to start the upgrade for a specific node or HA pair.

1. Validating the package and verifying the validation results

The following example shows how to validate a cluster software package. You must validate the package before the software upgrade. Set the `validate_only` field to `true` to start the validation. You can check for validation results in the `GET /cluster/software` endpoint.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?validate_only=true"
-H "accept: application/json" -H "Content-Type: application/hal+json" -d
'{ "version": "9.5.0" }'

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to validate the software cluster version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the `state` field of the job is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "PATCH /api/cluster/software",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

You can check for validation results in the GET /cluster/software endpoint. The following example shows how to check the validation warnings and errors after setting the `validate_only` field to `true`.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept: application/hal+json"

# The response:
{
  "version": "9.7.0",
  "validation_results": [
    {
      "update_check": "High Availability status",
      "status": "error",
      "issue": {
        "message": "Cluster HA is not configured in the cluster. Storage failover is not enabled on node \"node1\", \"node2\".",
      },
      "action": {
```



```

    "message": "Check cluster HA configuration. Check storage failover
status."
  },
  {
    "update_check": "Manual checks",
    "status": "warning",
    "issue" : {
      "message": "Manual validation checks need to be performed. Refer to
the Upgrade Advisor Plan or \"Performing manual checks before an automated
cluster upgrade\" section in the \"Clustered Data ONTAP Upgrade Express
Guide\" for the remaining validation checks that need to be performed
before update. Failing to do so can result in an update failure or an I/O
disruption."
    },
    "action": {
      "message": "Refer to the Upgrade Advisor Plan or \"Performing manual
checks before an automated cluster upgrade\" section in the \"Clustered
Data ONTAP Upgrade Express Guide\" for the remaining validation checks
that need to be performed before update."
    }
  }
],
"nodes": [
  {
    "node": "node1",
    "version": "9.7.0"
  },
  {
    "node": "node2",
    "version": "9.7.0"
  }
],
"state": "failed",
"elapsed_duration": 56,
"estimated_duration": 600,
"_links": {
  "self": {
    "href": "/api/cluster/software"
  }
}
}

```

2. Updating the cluster

The following example shows how to initiate a cluster software upgrade. You must validate the package before the software upgrade starts. Set the `skip_warnings` field to `true` to skip validation warnings and start the software package upgrade. You can specify the `stabilize_minutes` value between 1 to 60 minutes. Setting `stabilize_minutes` allows each node a specified amount of time to stabilize after a reboot; the default is 8 minutes.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?skip_warnings=true"
-H "accept: application/json" -H "Content-Type: application/hal+json" -d
'{ "version": "9.5.0" }'

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to update the software cluster version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the `state` field of the job is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "PATCH /api/cluster/software",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

You can check the update progress information in the GET /cluster/software endpoint. The following example shows how to check the progress of an update after setting the `skip_warnings` field to true.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept: application/hal+json"

# The response:
{
  "version": "9.7.0",
  "validation_results": [
    {
      "update_check": "Manual checks",
      "status": "warning",
      "issue": {
        "message": "Manual validation checks need to be performed. Refer to the Upgrade Advisor Plan or \"Performing manual checks before an automated cluster upgrade\" section in the \"Clustered Data ONTAP Upgrade Express Guide\" for the remaining validation checks that need to be performed"
      }
    }
  ]
}
```

before update. Failing to do so can result in an update failure or an I/O disruption."

```
    },
    "action": {
        "message": "Refer to the Upgrade Advisor Plan or \"Performing manual
checks before an automated cluster upgrade\" section in the \"Clustered
Data ONTAP Upgrade Express Guide\" for the remaining validation checks
that need to be performed before update."
    }
}
],
"nodes": [
    {
        "node": "node1",
        "version": "9.7.0"
    },
    {
        "node": "node2",
        "version": "9.7.0"
    }
],
"pending_version": "9.7.0",
"state": "in_progress",
"elapsed_duration": 63,
"estimated_duration": 5220,
"status_details": [
    {
        "name": "do-download-job",
        "status": "running",
        "issue": {
            "message": "Installing software image.",
            "code": 10551400
        },
        "start_time": "2019-01-14T23:12:14+05:30",
        "end_time": "2019-01-14T23:12:14+05:30",
        "node": {
            "name": "node1"
        }
    },
    {
        "name": "do-download-job",
        "status": "running",
        "issue": {
            "message": "Installing software image.",
            "code": 10551400
        }
    },

```

```

    "start_time": "2019-01-14T23:12:14+05:30",
    "end_time": "2019-01-14T23:12:14+05:30",
    "node": {
      "name": "node2"
    }
  },
  "update_details": [
    {
      "phase": "Data ONTAP updates",
      "status": "in-progress",
      "estimated_duration": 4620,
      "elapsed_duration": 10,
      "node": {
        "name": "node1"
      }
    },
    {
      "phase": "Data ONTAP updates",
      "status": "in-progress",
      "estimated_duration": 4620,
      "elapsed_duration": 10,
      "node": {
        "name": "node2"
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/cluster/software"
    }
  }
}

```

3. Pausing/resuming/canceling the upgrade

The following example shows how to pause an ongoing cluster software package upgrade. Set the `action` field to `pause`, `resume`, or `cancel` to pause, resume or cancel the upgrade respectively. Not all update operations support these actions. An update can only be resumed if it is in the paused state.

```
# The API:
/api/cluster/software

# The call:
curl -X PATCH "https://<mgmt_ip>/api/cluster/software?action=pause" -H
"accept: application/json" -H "Content-Type: application/hal+json" -d '{
"version": "9.5.0"}'

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to update the software cluster version returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the `state` field of the job is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "PATCH /api/cluster/software",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

You can check the progress of the upgrade in the GET /cluster/software endpoint. The following example shows how to check the progress of the pause upgrade state after setting the action field to pause.

```
# The API:
/api/cluster/software

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software" -H "accept: application/hal+json"

# The response:
{
  "version": "9.7.0",
  "validation_results": [
    {
      "update_check": "Manual checks",
      "status": "warning",
      "issue": {
        "message": "Manual validation checks need to be performed. Refer to the Upgrade Advisor Plan or \"Performing manual checks before an automated cluster upgrade\" section in the \"Clustered Data ONTAP Upgrade Express Guide\" for the remaining validation checks that need to be performed"
      }
    }
  ]
}
```

```

before update. Failing to do so can result in an update failure or an I/O
disruption."
    },
    "action": {
        "message": "Refer to the Upgrade Advisor Plan or \"Performing manual
checks before an automated cluster upgrade\" section in the \"Clustered
Data ONTAP Upgrade Express Guide\" for the remaining validation checks
that need to be performed before update."
    }
}
],
"nodes": [
    {
        "node": "node1",
        "version": "9.7.0"
    },
    {
        "node": "node2",
        "version": "9.7.0"
    }
],
"pending_version": "9.7.0",
"state": "pause_pending",
"elapsed_duration": 103,
"estimated_duration": 5220,
"status_details": [
    {
        "status": "in-progress",
        "issue": {
            "message": "Installing software image.",
            "code": 10551400
        },
        "start_time": "2019-01-08T02:54:36+05:30",
        "node": {
            "name": "node1"
        }
    },
    {
        "status": "in-progress",
        "issue": {
            "message": "Installing software image.",
            "code": 10551400
        },
        "start_time": "2019-01-08T02:54:36+05:30",
        "node": {
            "name": "node2"
        }
    }
]

```



```

    }
  }
],
"update_details": [
  {
    "phase": "Pre-update checks",
    "status": "completed",
    "estimated_duration": 600,
    "elapsed_duration": 54,
    "node": {
      "name": "node1"
    }
  },
  {
    "phase": "Data ONTAP updates",
    "status": "pause-pending",
    "estimated_duration": 4620,
    "elapsed_duration": 49,
    "node": {
      "name": "node2"
    }
  },
  {
    "phase": "Data ONTAP updates",
    "status": "pause-pending",
    "estimated_duration": 4620,
    "elapsed_duration": 49
  }
],
"_links": {
  "self": {
    "href": "/api/cluster/software"
  }
}
}

```

Downloading the software package

The following example shows how to download the software package from an HTTP or FTP server. Provide the `url`, `username`, and `password`, if required, to start the download of the software package to the cluster.

```
# The API:
/api/cluster/software/download

# The call:
curl -X POST "https://<mgmt-
ip>/api/cluster/software/download?return_timeout=0" -H "accept:
application/json" -H "Content-Type: application/hal+json" -d '{ "url":
"http://nbsweb.eng.btc.netapp.in/~suvadipd/99/image1.tgz", "username":
"admin", "password": "*****"}'

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to download the software package returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the job `state` field is set to success.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "POST /api/cluster/software/download",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

Checking the progress of the software package being downloaded from an HTTP or FTP server

The following example shows how to retrieve the progress status of the software package being downloaded from a HTTP or FTP server.

```
# The API:
/api/cluster/software/download

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software/download" -H "accept:
application/hal+json"

# The response:
{
  "state": "running",
  "message": "Package download in progress",
  "code": 10551760,
  "_links": {
    "self": {
      "href": "/api/cluster/software/download"
    }
  }
}
```

Uploading a software package

The following example shows how to upload a software package.

```
# The API:
/api/cluster/software/upload

# The call:
curl -ku username:password -F "file=@image.tgz" -X POST "https://<mgmt-
ip>/api/cluster/software/upload?return_timeout=0"

# The response:
{
  "job": {
    "uuid": "12db53fd-8326-11ea-91eb-005056bb16e5",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/12db53fd-8326-11ea-91eb-005056bb16e5"
      }
    }
  }
}
```

HTTPS error codes

The following is a list of possible error codes that can be returned during a package upload operation.

ONTAP Error Response codes

Error codes	Description
10551799	Internal error. Failed to check if file upload is enabled. Contact technical support for assistance.
10551800	File upload is disabled. Enable file upload by setting "ApacheUploadEnabled 1" in the web services configuration file or contact technical support for assistance.
10551802	Internal error. Access permissions restrict file upload. This is likely due to a bad web jail setup. Contact technical support for assistance.
10551803	Internal error. A read/write error occurred when uploading this file. Contact technical support for assistance.
10551804	An invalid argument was supplied to create a file handle. Try uploading the file again or contact technical support for assistance.
10551805	An unknown error occurred. Retry file upload operation again or contact technical support for assistance.
10551806	Internal error. There is not sufficient space in the file upload directory to upload this file. Contact technical support for assistance.
10551807	Internal error in JAIL setup. Contact technical support for assistance.
10551808	Internal error. Failed to write to file in the webjail directory. Contact technical support for assistance.
10551809	The request must only contain a single file. More than one file per request is not supported.

Retrieving cluster software packages information

The following example shows how to retrieve the ONTAP software packages in a cluster.

```
# The API:
/api/cluster/software/packages

# The call:
curl -X GET "https://<mgmt-
ip>/api/cluster/software/packages?return_records=true&return_timeout=15"
-H "accept: application/hal+json"

# The response:
{
  "records": [
    {
      "version": "9.7.0",
      "_links": {
        "self": {
          "href": "/api/cluster/software/packages/9.7.0"
        }
      }
    },
    {
      "version": "9.5.0",
      "_links": {
        "self": {
          "href": "/api/cluster/software/packages/9.5.0"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/cluster/software/packages"
    }
  }
}
```

The following example shows how to retrieve the details of a given cluster software package.

```
# The API:
/api/cluster/software/packages/{version}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software/packages/9.7.0" -H
"accept: application/hal+json"

# The response:
{
  "version": "9.7.0",
  "create_time": "2018-05-21T10:06:59+05:30",
  "_links": {
    "self": {
      "href": "/api/cluster/software/packages/9.7.0"
    }
  }
}
```

Deleting a cluster software package

The following example shows how to delete a package from the cluster. You need to provide the package version that you want to delete. The software package delete creates a job to perform the delete operation.

```
# The API:
/api/cluster/software/packages/{version}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/cluster/software/packages/9.6.0" -H
"accept: application/hal+json"

# The response:
{
  "job": {
    "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
      }
    }
  }
}
```

The call to delete the package returns the job UUID, including a HAL link to retrieve details about the job. The job object includes a `state` field and a message to indicate the progress of the job. When the job is complete and the application is fully created, the message indicates success and the job `state` field is set to `success`.

```
# The API:
/api/cluster/jobs/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc" -H "accept: application/hal+json"

# The response:
{
  "uuid": "f587d316-5feb-11e8-b0e0-005056956dfc",
  "description": "DELETE /api/cluster/software/packages/9.6.0",
  "state": "success",
  "message": "success",
  "code": 0,
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/f587d316-5feb-11e8-b0e0-005056956dfc"
    }
  }
}
```

HTTPS error codes

The following is a list of possible error codes that can be returned during a package delete operation.

ONTAP Error Response codes

Error codes	Description
10551315	Package store is empty
10551322	Error in retrieving package cleanup status
10551323	Error in cleaning up package information on a node
10551324	Error in cleaning up package information on both nodes
10551325	Package does not exist on the system
10551326	Error in deleting older package cleanup tasks
10551346	Package delete failed since a validation is in progress
10551347	Package delete failed since an update is in progress

Error codes	Description
10551367	A package synchronization is in progress
10551388	Package delete operation timed out

Retrieving software installation history information

The following example shows how to:

- retrieve the software package installation history information.
- display specific node level software installation history information.
- provide all the attributes by default in response when the self referential link is not present.

```
# The API:
/api/cluster/software/history

# The call:
curl -X GET "https://<mgmt-ip>/api/cluster/software/history" -H "accept:
application/hal+json"

# The response:
{
  "node": {
    "uuid": "58cd3a2b-af63-11e8-8b0d-0050568e7279",
    "name": "sti70-vsim-ucsl65n",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/58cd3a2b-af63-11e8-8b0d-0050568e7279"
      }
    }
  },
  "start_time": "2018-09-03T16:18:46+05:30",
  "state": "successful",
  "from_version": "9.4.0",
  "to_version": "9.5.0",
  "end_time": "2018-05-21T10:14:51+05:30"
}
```

Retrieve the cluster software profile

GET /cluster/software

Introduced In: 9.6

Retrieves the software profile of a cluster.

Related ONTAP commands

- `cluster image show`
- `cluster image show-update-progress`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False
Specify the fields to return.	return_timeout	integer	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
action	string	User triggered action to apply to the install operation
elapsed_duration	integer	Elapsed time during the upgrade or validation operation

Name	Type	Description
estimated_duration	integer	Estimated time remaining until completion of the upgrade or validation operation.
metrocluster	metrocluster	
nodes	array[software_node_reference]	List of nodes and active versions.
pending_version	string	Version being installed on the system. <ul style="list-style-type: none"> • example: ONTAP_X_1 • readOnly: 1 • Introduced in: 9.6
state	string	Operational state of the upgrade
status_details	array[software_status_details_reference]	Display status details.
update_details	array[software_update_details_reference]	Display update progress details.
validation_results	array[software_validation_reference]	List of validation warnings, errors, and advice.
version	string	Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation. <ul style="list-style-type: none"> • example: ONTAP_X • Introduced in: 9.6

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "action": "pause",
  "elapsed_duration": "2140",
  "estimated_duration": "5220",
  "metrocluster": {
    "clusters": [
      {
        "elapsed_duration": "2140",
        "estimated_duration": "3480",
        "name": "cluster_A",
        "state": "in_progress"
      }
    ],
    "progress_details": {
      "message": "Switchover in progress"
    },
    "progress_summary": {
      "message": "MetroCluster updated successfully."
    }
  },
  "nodes": [
    {
      "name": "node1",
      "version": "ONTAP_X"
    }
  ],
  "pending_version": "ONTAP_X_1",
  "state": "completed",
  "status_details": [
    {
      "action": {
        "message": "string"
      },
      "end_time": "2019-02-02T19:00:00Z",
      "issue": {
        "code": "10551399",
        "message": "Image update complete"
      },
      "name": "initialize",

```

```

    "node": {
      "name": "node1"
    },
    "start_time": "2019-02-02T19:00:00Z",
    "state": "failed"
  }
],
"update_details": [
  {
    "elapsed_duration": "2100",
    "estimated_duration": "4620",
    "node": {
      "name": "node1"
    },
    "phase": "Pre-update checks",
    "state": "failed"
  }
],
"validation_results": [
  {
    "action": {
      "message": "string"
    },
    "issue": {
      "message": "Validation error: Cluster HA is not configured in
the cluster"
    },
    "status": "warning",
    "update_check": "nfs_mounts"
  }
],
"version": "ONTAP_X"
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

software_mcc_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration of update time (in seconds) of MetroCluster.
estimated_duration	integer	Estimated duration of update time (in seconds) of MetroCluster.
name	string	Name of the site in MetroCluster.
state		Upgrade state of MetroCluster.

progress_details

Name	Type	Description
message	string	MetroCluster update progress details.

progress_summary

Name	Type	Description
message	string	MetroCluster update progress summary.

metrocluster

Name	Type	Description
clusters	array[software_mcc_reference]	List of MetroCluster sites, statuses, and active ONTAP versions. <ul style="list-style-type: none"> • readOnly: 1 • Introduced in: 9.6
progress_details	progress_details	
progress_summary	progress_summary	

software_node_reference

Name	Type	Description
name	string	Name of the node.
version	string	ONTAP version of the node. <ul style="list-style-type: none"> • example: ONTAP_X • readOnly: 1 • Introduced in: 9.6

action

Name	Type	Description
code	integer	Error code corresponding the status error
message	string	Corrective action to be taken to resolve the status error.

issue

Name	Type	Description
code	integer	Error code corresponding to update status
message	string	Update status details

node

Name	Type	Description
name	string	Name of the node to be retrieved for status details.

software_status_details_reference

Name	Type	Description
action	action	
end_time	string	End time for each status phase.
issue	issue	
name	string	Name of the phase to be retrieved for status details.
node	node	
start_time	string	Start time for each status phase.
state	string	Status of the phase

node

Name	Type	Description
name	string	Name of the node to be retrieved for update details.

software_update_details_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration for each update phase
estimated_duration	integer	Estimated duration for each update phase
node	node	
phase	string	Phase details
state	string	State of the update phase

action

Name	Type	Description
message	string	Specifies the corrective action to be taken to resolve a validation error

issue

Name	Type	Description
message	string	Details of the error or warning encountered by the update checks

software_validation_reference

Name	Type	Description
action	action	
issue	issue	
status	string	Status of this update check.
update_check	string	Name of the update check to be validated.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array [error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the cluster software version

PATCH /cluster/software

Introduced In: 9.6

Updates the cluster software version.

Important note:

- Setting 'version' triggers the package installation.
- To validate the package for installation but not perform the installation, use the `validate_only` field on the request.

Required properties

- `version` - Software version to be installed on the cluster.

Recommended optional parameters

- `validate_only` - Required to validate a software package before an upgrade.
- `skip_warnings` - Used to skip validation warnings when starting a software upgrade.
- `action` - Used to pause, resume, or cancel an ongoing software upgrade.
- `stabilize_minutes` - Specifies a custom value between 1 to 60 minutes that allows each node a specified amount of time to stabilize after a reboot; the default is 8 minutes.

Related ONTAP commands

- `cluster image validate`
- `cluster image update`
- `cluster image pause-update`
- `cluster image resume-update`
- `cluster image cancel-update`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
skip_warnings	boolean	query	False	Ignore warnings and proceed with the install.

Name	Type	In	Required	Description
action	string	query	False	<p>Requests an upgrade to pause, resume, or cancel. Note that not all upgrades support these actions. An upgrade can only be resumed if it is in the paused state. When a request to cancel an upgrade is successful, the upgrade state changes to either success or failure.</p> <ul style="list-style-type: none"> enum: ["pause", "resume", "cancel"]
stabilize_minutes	integer	query	False	<p>Sets a custom value between 1 to 60 minutes for the upgrade, allowing each node a specified amount of time to stabilize after a reboot.</p> <ul style="list-style-type: none"> Introduced in: 9.8
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
validate_only	boolean	query	False	<p>Validate the operation and its parameters, without actually performing the operation.</p>

Request Body

Name	Type	Description
action	string	User triggered action to apply to the install operation
elapsed_duration	integer	Elapsed time during the upgrade or validation operation

Name	Type	Description
estimated_duration	integer	Estimated time remaining until completion of the upgrade or validation operation.
metrocluster	metrocluster	
nodes	array[software_node_reference]	List of nodes and active versions.
pending_version	string	Version being installed on the system. <ul style="list-style-type: none"> • example: ONTAP_X_1 • readOnly: 1 • Introduced in: 9.6
state	string	Operational state of the upgrade
status_details	array[software_status_details_reference]	Display status details.
update_details	array[software_update_details_reference]	Display update progress details.
validation_results	array[software_validation_reference]	List of validation warnings, errors, and advice.
version	string	Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation. <ul style="list-style-type: none"> • example: ONTAP_X • Introduced in: 9.6

Example request

```
{
  "action": "pause",
  "elapsed_duration": "2140",
  "estimated_duration": "5220",
  "metrocluster": {
    "clusters": [
      {
        "elapsed_duration": "2140",
        "estimated_duration": "3480",
        "name": "cluster_A",
        "state": "in_progress"
      }
    ],
    "progress_details": {
      "message": "Switchover in progress"
    },
    "progress_summary": {
      "message": "MetroCluster updated successfully."
    }
  },
  "nodes": [
    {
      "name": "node1",
      "version": "ONTAP_X"
    }
  ],
  "pending_version": "ONTAP_X_1",
  "state": "completed",
  "status_details": [
    {
      "action": {
        "message": "string"
      },
      "end_time": "2019-02-02T19:00:00Z",
      "issue": {
        "code": "10551399",
        "message": "Image update complete"
      },
      "name": "initialize",
      "node": {
        "name": "node1"
      },
      "start_time": "2019-02-02T19:00:00Z",
      "state": "failed"
    }
  ]
}
```

```

    }
  ],
  "update_details": [
    {
      "elapsed_duration": "2100",
      "estimated_duration": "4620",
      "node": {
        "name": "node1"
      },
      "phase": "Pre-update checks",
      "state": "failed"
    }
  ],
  "validation_results": [
    {
      "action": {
        "message": "string"
      },
      "issue": {
        "message": "Validation error: Cluster HA is not configured in
the cluster"
      },
      "status": "warning",
      "update_check": "nfs_mounts"
    }
  ],
  "version": "ONTAP_X"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{  
  "job": {  
    "uuid": "string"  
  }  
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

software_mcc_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration of update time (in seconds) of MetroCluster.
estimated_duration	integer	Estimated duration of update time (in seconds) of MetroCluster.
name	string	Name of the site in MetroCluster.
state		Upgrade state of MetroCluster.

progress_details

Name	Type	Description
message	string	MetroCluster update progress details.

progress_summary

Name	Type	Description
message	string	MetroCluster update progress summary.

metrocluster

Name	Type	Description
clusters	array[software_mcc_reference]	List of MetroCluster sites, statuses, and active ONTAP versions. <ul style="list-style-type: none">readOnly: 1Introduced in: 9.6
progress_details	progress_details	

Name	Type	Description
progress_summary	progress_summary	

software_node_reference

Name	Type	Description
name	string	Name of the node.
version	string	ONTAP version of the node. <ul style="list-style-type: none"> • example: ONTAP_X • readOnly: 1 • Introduced in: 9.6

action

Name	Type	Description
code	integer	Error code corresponding the status error
message	string	Corrective action to be taken to resolve the status error.

issue

Name	Type	Description
code	integer	Error code corresponding to update status
message	string	Update status details

node

Name	Type	Description
name	string	Name of the node to be retrieved for status details.

software_status_details_reference

Name	Type	Description
action	action	

Name	Type	Description
end_time	string	End time for each status phase.
issue	issue	
name	string	Name of the phase to be retrieved for status details.
node	node	
start_time	string	Start time for each status phase.
state	string	Status of the phase

node

Name	Type	Description
name	string	Name of the node to be retrieved for update details.

software_update_details_reference

Name	Type	Description
elapsed_duration	integer	Elapsed duration for each update phase
estimated_duration	integer	Estimated duration for each update phase
node	node	
phase	string	Phase details
state	string	State of the update phase

action

Name	Type	Description
message	string	Specifies the corrective action to be taken to resolve a validation error

issue

Name	Type	Description
message	string	Details of the error or warning encountered by the update checks

software_validation_reference

Name	Type	Description
action	action	
issue	issue	
status	string	Status of this update check.
update_check	string	Name of the update check to be validated.

software_reference

Name	Type	Description
action	string	User triggered action to apply to the install operation
elapsed_duration	integer	Elapsed time during the upgrade or validation operation
estimated_duration	integer	Estimated time remaining until completion of the upgrade or validation operation.
metrocluster	metrocluster	
nodes	array[software_node_reference]	List of nodes and active versions.
pending_version	string	Version being installed on the system. <ul style="list-style-type: none"> • example: ONTAP_X_1 • readOnly: 1 • Introduced in: 9.6
state	string	Operational state of the upgrade
status_details	array[software_status_details_reference]	Display status details.
update_details	array[software_update_details_reference]	Display update progress details.

Name	Type	Description
validation_results	array[software_validation_reference]	List of validation warnings, errors, and advice.
version	string	<p>Version of ONTAP installed and currently active on the system. During PATCH, using the 'validate_only' parameter on the request executes pre-checks, but does not perform the full installation.</p> <ul style="list-style-type: none"> • example: ONTAP_X • Introduced in: 9.6

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the software or firmware download status

GET /cluster/software/download

Introduced In: 9.7

Retrieves the software download status.

Related ONTAP commands

- `cluster image package check-download-progress`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False
Specify the fields to return.	return_timeout	integer	query	False

Response

Status: 200, Ok

Name	Type	Description
code	integer	Code corresponds to download message

Name	Type	Description
message	string	Download progress details
state	string	Download status of the package

Example response

```
{
  "code": "10551496",
  "message": "Package download in progress",
  "state": "success"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Download a software or firmware package

POST `/cluster/software/download`

Introduced In: 9.6

Downloads a software package from the server.

Required properties

- `url` - URL location of the software package

Recommended optional parameters

- `username` - Username of HTTPS/FTP server
- `password` - Password of HTTPS/FTP server

Related ONTAP commands

- `cluster image package get`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none">• Default value:
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Request Body

Name	Type	Description
password	string	Password for download

Name	Type	Description
url	string	HTTP or FTP URL of the package through a server
username	string	Username for download

Example request

```
{
  "password": "admin_password",
  "url": "http://server/package",
  "username": "admin"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

software_package_download

Name	Type	Description
password	string	Password for download
url	string	HTTP or FTP URL of the package through a server
username	string	Username for download

href

Name	Type	Description
href	string	

_links

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Retrieve the software installation request history details

GET /cluster/software/history

Introduced In: 9.6

Retrieves the history details for software installation requests.

Related ONTAP commands

- `cluster image show-update-history`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
end_time	string	query	False	Filter by end_time <ul style="list-style-type: none"> • Introduced in: 9.7
node.name	string	query	False	Filter by node.name <ul style="list-style-type: none"> • Introduced in: 9.7
node.uuid	string	query	False	Filter by node.uuid <ul style="list-style-type: none"> • Introduced in: 9.7
to_version	string	query	False	Filter by to_version <ul style="list-style-type: none"> • Introduced in: 9.7
start_time	string	query	False	Filter by start_time <ul style="list-style-type: none"> • Introduced in: 9.7

Name	Type	In	Required	Description
state	string	query	False	Filter by state <ul style="list-style-type: none"> Introduced in: 9.7
from_version	string	query	False	Filter by from_version <ul style="list-style-type: none"> Introduced in: 9.7
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none"> Default value: 1 Max value: 120 Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	
records	array[software_history]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "end_time": "2019-02-02T20:00:00Z",
      "from_version": "ONTAP_X1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "start_time": "2019-02-02T19:00:00Z",
      "state": "successful",
      "to_version": "ONTAP_X2"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

software_history

Name	Type	Description
end_time	string	Completion time of this installation request.
from_version	string	Previous version of node <ul style="list-style-type: none">• example: ONTAP_X1• readOnly: 1• Introduced in: 9.7
node	node	
start_time	string	Start time of this installation request.
state	string	Status of this installation request.

Name	Type	Description
to_version	string	Updated version of node <ul style="list-style-type: none"> • example: ONTAP_X2 • readOnly: 1 • Introduced in: 9.7

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve cluster software packages

GET `/cluster/software/packages`

Introduced In: 9.6

Retrieves the software packages for a cluster.

Related ONTAP commands

- `cluster image package show-repository`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
version	string	query	False	Filter by version <ul style="list-style-type: none"> Introduced in: 9.7
create_time	string	query	False	Filter by create_time <ul style="list-style-type: none"> Introduced in: 9.7
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none"> Default value: 1 Max value: 120 Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	
records	array[software_package]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "create_time": "2019-02-04T19:00:00Z",
      "version": "ONTAP_X"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

software_package

Name	Type	Description
_links	_links	
create_time	string	Indicates when this package was loaded
version	string	Version of this package <ul style="list-style-type: none">• example: ONTAP_X• readOnly: 1• Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a software package from the cluster

DELETE /cluster/software/packages/{version}

Introduced In: 9.6

Deletes a software package from the cluster. The delete operation fails if the package is currently installed.

Related ONTAP commands

- `cluster image package delete`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
version	string	path	True	

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Error

Status: Default

ONTAP error response codes

Error codes	Description
10551315	Package store is empty

Error codes	Description
10551322	Error in retrieving package cleanup status
10551323	Error in cleaning up package information on a node
10551324	Error in cleaning up package information on multiple nodes
10551325	Package does not exist on the system
10551326	Error in deleting older package cleanup tasks. Clean up images from the store and retry
10551346	Package delete failed since a validation is in progress
10551347	Package delete failed since an update is in progress
10551367	A package synchronization is in progress
10551388	Package delete operation timed out

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the software package information

GET /cluster/software/packages/{version}

Introduced In: 9.6

Retrieves the software package information.

Related ONTAP commands

- `cluster image package show-repository`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
version	string	path	True	
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
create_time	string	Indicates when this package was loaded
version	string	Version of this package <ul style="list-style-type: none">• example: ONTAP_X• readOnly: 1• Introduced in: 9.6

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "create_time": "2019-02-04T19:00:00Z",
  "version": "ONTAP_X"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Upload a software or firmware package located on the local file system

POST /cluster/software/upload

Introduced In: 9.8

Uploads a software package file located on the local filesystem.

Related ONTAP commands

- `cluster image package get`

Learn more

- [DOC /cluster/software](#)

Parameters

Name	Type	In	Required	Description
file	file	formData	False	Info specification
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none">• Default value:
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage

Private manage event-remediations

Private manage event-remediations endpoint overview

Overview

Event remediations is the initial implementation of ONTAP self healing proof of concept.

This endpoint supports GET calls. GET is used to retrieve details about the event.

Event remediation and management action collection get

learn more

- [\[doc /private/manage/event-remediations\]\(#docs-manage-private_manage_event-remediations\)](#)

GET /private/manage/event-remediations

Introduced In: 9.8

Event remediation and management action collection GET

Learn more

- [DOC /private/manage/event-remediations](#)

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[event_remediations]	List of event remediations

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "3",
  "records": [
    {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
      },
      "completion_time": "string",
      "description": "Disable Telnet",
      "event_remediation_action": "string",
      "event_remediation_action_type": "string",
      "event_remediation_state": "string",
      "event_type_name": "cluster_telnet_enabled",
      "id": "198",
      "solution": "Continue to disable telnet on the selected cluster using the 'security protocol modify -application telnet -enabled false' ontap cli command.",
      "source": {
        "object_id": "5666",
        "object_type": "inventory.ontap.fas.Cluster"
      },
      "source_full_name": "shrey-vsml",
      "source_resource_type": "cluster",
      "start_time": "string"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

cluster

Name	Type	Description
name	string	
uuid	string	

source

Name	Type	Description
object_id	integer	
object_type	string	

event_remediations

Name	Type	Description
cluster	cluster	
completion_time	string	Completion time of the event
description	string	Description for the event
event_remediation_action	string	Event remediation action
event_remediation_action_type	string	Type of remediation action
event_remediation_state	string	State of the event remediation
event_type_name	string	Type of the event
id	integer	Event Identifier

Name	Type	Description
solution	string	Corrective action for the event
source	source	
source_full_name	string	Source of the event
source_resource_type	string	Type of source of the event
start_time	string	Start time of the event

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Event remediation and management action

learn more

- [\[doc /private/manage/event-remediations\]\(#docs-manage-private_manage_event-remediations\)](#)

GET /private/manage/event-remediations/{id}

Introduced In: 9.8

Event remediation and management action

Learn more

- [DOC /private/manage/event-remediations](#)

Parameters

Name	Type	In	Required	Description
id	integer	path	True	• Default value: 1
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
cluster	cluster	
completion_time	string	Completion time of the event
description	string	Description for the event
event_remediation_action	string	Event remediation action
event_remediation_action_type	string	Type of remediation action
event_remediation_state	string	State of the event remediation
event_type_name	string	Type of the event
id	integer	Event Identifier
solution	string	Corrective action for the event
source	source	
source_full_name	string	Source of the event
source_resource_type	string	Type of source of the event
start_time	string	Start time of the event

Example response

```
{
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "completion_time": "string",
  "description": "Disable Telnet",
  "event_remediation_action": "string",
  "event_remediation_action_type": "string",
  "event_remediation_state": "string",
  "event_type_name": "cluster_telnet_enabled",
  "id": "198",
  "solution": "Continue to disable telnet on the selected cluster using the 'security protocol modify -application telnet -enabled false' ontap cli command.",
  "source": {
    "object_id": "5666",
    "object_type": "inventory.ontap.fas.Cluster"
  },
  "source_full_name": "shrey-vsimpl",
  "source_resource_type": "cluster",
  "start_time": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

cluster

Name	Type	Description
name	string	
uuid	string	

source

Name	Type	Description
object_id	integer	
object_type	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Perform an event management action

learn more

- [\[doc /private/manage/event-remediations\]\(#docs-manage-private_manage_event-remediations\)](#)

PATCH /private/manage/event-remediations/{id}

Introduced In: 9.8

Perform an event management action.

Learn more

- [DOC /private/manage/event-remediations](#)

Parameters

Name	Type	In	Required	Description
id	integer	path	True	<ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Request Body

Name	Type	Description
cluster	cluster	
completion_time	string	Completion time of the event

Name	Type	Description
description	string	Description for the event
event_remediation_action	string	Event remediation action
event_remediation_action_type	string	Type of remediation action
event_remediation_state	string	State of the event remediation
event_type_name	string	Type of the event
id	integer	Event Identifier
solution	string	Corrective action for the event
source	source	
source_full_name	string	Source of the event
source_resource_type	string	Type of source of the event
start_time	string	Start time of the event

Example request

```
{
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "completion_time": "string",
  "description": "Disable Telnet",
  "event_remediation_action": "string",
  "event_remediation_action_type": "string",
  "event_remediation_state": "string",
  "event_type_name": "cluster_telnet_enabled",
  "id": "198",
  "solution": "Continue to disable telnet on the selected cluster using the 'security protocol modify -application telnet -enabled false' ontap cli command.",
  "source": {
    "object_id": "5666",
    "object_type": "inventory.ontap.fas.Cluster"
  },
  "source_full_name": "shrey-vsml",
  "source_resource_type": "cluster",
  "start_time": "string"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Ontap error response codes

Definitions

See Definitions

cluster

Name	Type	Description
name	string	
uuid	string	

source

Name	Type	Description
object_id	integer	
object_type	string	

event_remediations

Name	Type	Description
cluster	cluster	
completion_time	string	Completion time of the event
description	string	Description for the event
event_remediation_action	string	Event remediation action
event_remediation_action_type	string	Type of remediation action
event_remediation_state	string	State of the event remediation
event_type_name	string	Type of the event
id	integer	Event Identifier
solution	string	Corrective action for the event
source	source	
source_full_name	string	Source of the event
source_resource_type	string	Type of source of the event
start_time	string	Start time of the event

href

Name	Type	Description
href	string	

_links

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

NAS

NAS overview

Overview

These APIs allow you to complete various tasks, including:

- Creating an NFS server for an SVM
- Managing an NFS configuration of an SVM
- Viewing and updating the NFS configuration of an SVM
- Configuring export policies and rules for an SVM

- Managing export policies and rules for an SVM

APIs

NFS

The NFS APIs enable you to create and configure NFS settings for an SVM. You can delete or update NFS configurations, and you can also disable or enable different NFS features as needed.

Exports

The export APIs allow you to create and manage export policies for an SVM that enable an administrator to restrict access to volumes for clients that match specific IP addresses and specific authentication types. Export APIs are also used to create export rules for an export policy. The APIs allow each rule to specify the number of mask bits in the client IP address that must be matched for that rule to apply to a particular client request. The APIs also allow each export rule to specify the authentication types that are required for both read-only and read-write operations.

Kerberos

Kerberos is a protocol designed to provide strong authentication for users and hosts within a client/server environment. The basis of the protocol is a shared, secret-key cryptology system. (Kerberos uses shared-key encryption to ensure the confidentiality of the data. It also uses hashing techniques to ensure the integrity of the data (so that no one can modify the data unless allowed to do so). With the NetApp multiprotocol storage platform, through which clients based on UNIX or Windows can access data using CIFS or NFS, it is crucial to provide the ability to use standard network services for authentication and for identity storage.

To configure an ONTAP system to use Kerberos for NFS, Kerberos must be enabled on a data LIF in the SVM that owns the NFS server. A Kerberos realm needs to be created before enabling Kerberos on a data LIF. (The Kerberos realm is needed so that the cluster knows how to format Kerberos ticket requests.) The Kerberos APIs allow you to define, create, modify, and delete realms for the SVM. The APIs also allow you to enable/disable Kerberos on a data LIF and update the Kerberos interface configuration for a particular data LIF in the SVM.

Private protocols audit audit-log-redirect

Private protocols audit audit-log-redirect endpoint overview

Overview

Use this API to specify a dedicated SVM for all NAS auditing events log to reside in.

Examples

Creating an Audit Log Redirect configuration.

To create an Audit Log Redirect configuration.

```
# The API:
POST /api/private/protocols/audit/audit-log-redirect/

# The call:
curl -X POST "https://<mgmt-ip>/api/private/protocols/audit/audit-log-redirect?return_timeout=5&return_records=false" -H "accept: application/json" -H "authorization: Basic YWRtaW46bmV0YXBwMSE=" -H "Content-Type: application/json" -d "{ \"svm\": { \"uuid\": \"02c9e252-41be-11e9-81d5-00a0986138f7\" } }"

# The response:
{ }
```

Retrieving an Audit Log Redirect configuration in the cluster.

To retrieve an Audit Log Redirect configuration.

```
# The API:
GET /api/private/protocols/audit/audit-log-redirect/

# The call:
curl -X GET "https://<mgmt-ip>/api/private/protocols/audit/audit-log-redirect?return_records=true" -H "accept: application/json" -H "authorization: Basic YWRtaW46bmV0YXBwMSE="

# The response:
{
  "svm": {
    "uuid": "24870c49-8a73-11ea-ad05-005056827898",
    "name": "vs2"
  }
}
```

Updating Audit Log Redirect configuration in the cluster.

To modify an existing Audit Log Redirect configuration.

```
# The API:
PATCH /api/private/protocols/audit/audit-log-redirect/

# The call:
curl -X PATCH "https://<mgmt-ip>/api/private/protocols/audit/audit-log-redirect?return_records=true" -H "accept: application/json" -H "authorization: Basic YWRtaW46bmV0YXBwMSE=" -H "Content-Type: application/json" -d "{  \"svm\": {    \"name\": \"svm1\",    \"uuid\": \"02c9e252-41be-11e9-81d5-00a0986138f7\"  }}"
```

Deleting an Audit Log Redirect configuration in the cluster.

To delete an existing Audit Log Redirect configuration.

```
# The API:
DELETE /api/private/protocols/audit/audit-log-redirect/

# The call:
curl -X DELETE "https://<mgmt-ip>/api/private/protocols/audit/audit-log-redirect" -H "accept: application/json" -H "authorization: Basic YWRtaW46bmV0YXBwMSE="
```

Deletes an audit log redirect configuration

related ontap commands

- `vserver audit audit-log-redirect delete`
learn more
- `[doc /private/protocols/audit/audit-log-redirect](#docs-nas-private_protocols_audit_audit-log-redirect)`

```
DELETE /private/protocols/audit/audit-log-redirect
```

Introduced In: 9.8

Deletes an Audit Log Redirect configuration.

Related ONTAP commands

- `vserver audit audit-log-redirect delete`

Learn more

- [DOC /private/protocols/audit/audit-log-redirect](#)

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
9699482	Failed to delete Audit Log Redirect configuration.
9699484	Failed to delete Audit Log Redirect configuration because all SVMs audit configuration must be deleted first.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieves audit log redirect configuration

related ontap commands

- `vserver audit audit-log-redirect show`
[# learn more](#)
- `[doc /private/protocols/audit/audit-log-redirect](#docs-nas-private_protocols_audit_audit-log-redirect)`

GET /private/protocols/audit/audit-log-redirect

Introduced In: 9.8

Retrieves Audit Log Redirect configuration.

Related ONTAP commands

- `vserver audit audit-log-redirect show`

Learn more

- [DOC /private/protocols/audit/audit-log-redirect](#)

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
svm	svm	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Updates an audit log redirect configuration

required properties

- `svmuuid` or `svmname` - existing svm to which nas audit events logs will be redirected to
related ontap commands
- `vserver audit audit-log-redirect modify`
learn more
- `[doc /private/protocols/audit/audit-log-redirect](#docs-nas-private_protocols_audit_audit-log-redirect)`

PATCH `/private/protocols/audit/audit-log-redirect`

Introduced In: 9.8

Updates an Audit Log Redirect configuration.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM to which NAS audit events logs will be redirected to.

Related ONTAP commands

- `vserver audit audit-log-redirect modify`

Learn more

- [DOC /private/protocols/audit/audit-log-redirect](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
svm	svm	

Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
9699478	Audit Log Redirect configuration already exists.
9699480	Failed to modify Audit Log Redirect configuration.
9699481	Failed to modify Audit Log Redirect configuration because it does not exist.
9699483	Audit Log Redirect configuration not available.
9699484	Failed to modify Audit Log Redirect configuration because all SVMs audit configuration must be deleted first.
9699486	Audit Log Redirect feature is not available in maintenance mode.
9699487	Node is not part of a cluster.
9699488	SVM does not exist.
9699489	Audit Log Redirect configuration only supported on Data SVM.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

audit_log_redirect

Name	Type	Description
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Creates an audit log redirect configuration

required properties

- `svmuuid` or `svmname` - existing svm to which nas audit events logs will be redirected to
related ontap commands
- `vserver audit audit-log-redirect create`
learn more
- `[doc /private/protocols/audit/audit-log-redirect](#docs-nas-private_protocols_audit_audit-log-redirect)`

POST /private/protocols/audit/audit-log-redirect

Introduced In: 9.8

Creates an Audit Log Redirect configuration.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM to which NAS audit events logs will be redirected to.

Related ONTAP commands

- `vserver audit audit-log-redirect create`

Learn more

- [DOC /private/protocols/audit/audit-log-redirect](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
svm	svm	

Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 201, Created

Name	Type	Description
svm	svm	

Example response

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
9699478	Audit Log Redirect configuration already exists.
9699479	Failed to create Audit Log Redirect configuration.
9699483	Audit Log Redirect configuration not available.
9699484	Failed to modify Audit Log Redirect configuration because all SVMs audit configuration must be deleted first.

Error Code	Description
9699486	Audit Log Redirect feature is not available in maintenance mode.
9699487	Node is not part of a cluster.
9699488	SVM does not exist.
9699489	Audit Log Redirect configuration only supported on Data SVM.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

audit_log_redirect

Name	Type	Description
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage NAS audit configurations

Protocols audit endpoint overview

Overview

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on storage virtual machines (SVMs). This helps you track potential security problems and provides evidence of any security breaches.

Examples

Creating an audit entry with log rotation size and log retention count

To create an audit entry with log rotation size and log retention count, use the following API. Note the *return_records=true* query parameter is used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/audit/

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/audit" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
true, \"events\": { \"authorization_policy\": false, \"cap_staging\":
false, \"cifs_logon_logoff\": true, \"file_operations\": true,
\"file_share\": false, \"security_group\": false, \"user_account\": false
}, \"log\": { \"format\": \"evtx\", \"retention\": { \"count\": 10 },
\"rotation\": { \"size\": 2048000 }}, \"log_path\": \"/\", \"svm\": {
\"name\": \"vs1\", \"uuid\": \"ec650e97-156e-11e9-abcb-005056bbd0bf\" }}"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
        "name": "vs1"
      },
      "enabled": true,
      "events": {
        "authorization_policy": false,
        "cap_staging": false,
        "cifs_logon_logoff": true,
        "file_operations": true,
        "file_share": false,
        "security_group": false,
        "user_account": false
      }
    }
  ]
}
```



```

    },
    "log": {
        "format": "evtx",
        "rotation": {
            "size": 2048000
        },
        "retention": {
            "count": 10,
            "duration": "0s"
        }
    },
    "log_path": "/"
}
],
"num_records": 1
}

```

Creating an audit entry with log rotation schedule and log retention duration

To create an audit entry with log rotation schedule and log retention duration, use the following API. Note that the *return_records=true* query parameter is used to obtain the newly created entry in the response.

```

# The API:
POST /api/protocols/audit/

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/audit" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
false, \"events\": { \"authorization_policy\": false, \"cap_staging\":
false, \"cifs_logon_logoff\": true, \"file_operations\": true,
\"file_share\": false, \"security_group\": false, \"user_account\": false
}, \"log\": { \"format\": \"xml\", \"retention\": { \"duration\":
\"P4DT12H30M5S\" }, \"rotation\": { \"schedule\": { \"days\": [1, 5, 10,
15], \"hours\": [0, 1, 6, 12, 18, 23], \"minutes\": [10, 15, 30, 45, 59],
\"months\": [0], \"weekdays\": [0, 2, 5] } } }, \"log_path\": \"/\",
\"svm\": { \"name\": \"vs3\", \"uuid\": \"a8d64674-13fc-11e9-87b1-
005056a7ae7e\" } }"

# The response:
{
  "records": [
    {
      "svm": {

```

```

    "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
    "name": "vs3"
  },
  "enabled": true,
  "events": {
    "authorization_policy": false,
    "cap_staging": false,
    "cifs_logon_logoff": true,
    "file_operations": true,
    "file_share": false,
    "security_group": false,
    "user_account": false
  },
  "log": {
    "format": "xml",
    "rotation": {
      "schedule": {
        "minutes": [
          10,
          15,
          30,
          45,
          59
        ],
        "hours": [
          0,
          1,
          6,
          12,
          18,
          23
        ],
        "weekdays": [
          0,
          2,
          5
        ],
        "days": [
          1,
          5,
          10,
          15
        ],
        "months": [
          0
        ]
      ]
    }
  }
}

```

```

    }
  },
  "retention": {
    "count": 0,
    "duration": "P4DT12H30M5S"
  }
},
"log_path": "/"
}
],
"num_records": 1
}

```

Retrieving an audit configuration for all SVMs in the cluster

```

# The API:
GET /api/protocols/audit/

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/audit?fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
        "name": "vs1"
      },
      "enabled": true,
      "events": {
        "authorization_policy": false,
        "cap_staging": false,
        "cifs_logon_logoff": true,
        "file_operations": true,
        "file_share": false,
        "security_group": false,
        "user_account": false
      },
      "log": {

```

```

    "format": "evtx",
    "rotation": {
        "size": 2048000
    },
    "retention": {
        "count": 10,
        "duration": "0s"
    }
},
"log_path": "/"
},
{
    "svm": {
        "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
        "name": "vs3"
    },
    "enabled": true,
    "events": {
        "authorization_policy": false,
        "cap_staging": false,
        "cifs_logon_logoff": true,
        "file_operations": true,
        "file_share": false,
        "security_group": false,
        "user_account": false
    },
    "log": {
        "format": "xml",
        "rotation": {
            "schedule": {
                "minutes": [
                    10,
                    15,
                    30,
                    45,
                    59
                ],
                "hours": [
                    0,
                    1,
                    6,
                    12,
                    18,
                    23
                ]
            },
            "weekdays": [

```

```
        0,  
        2,  
        5  
    ],  
    "days": [  
        1,  
        5,  
        10,  
        15  
    ],  
    "months": [  
        0  
    ]  
    }  
},  
"retention": {  
    "count": 0,  
    "duration": "P4DT12H30M5S"  
}  
},  
"log_path": "/"  
}  
],  
"num_records": 2  
}
```

Retrieving specific entries with event list as cifs-logon-logout, file-ops = true for an SVM

The configuration returned is identified by the events in the list of audit configurations for an SVM.

```
# The API:
GET /api/protocols/audit/

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/audit?events.file_operations=true&events.cifs_logon_logoff=true&return_records=true&return_timeout=15" -H "accept:
application/json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
        "name": "vs1"
      },
      "events": {
        "cifs_logon_logoff": true,
        "file_operations": true
      }
    },
    {
      "svm": {
        "uuid": "a8d64674-13fc-11e9-87b1-005056a7ae7e",
        "name": "vs3"
      },
      "events": {
        "cifs_logon_logoff": true,
        "file_operations": true
      }
    }
  ],
  "num_records": 2
}
```

Retrieving a specific audit configuration for an SVM

The configuration returned is identified by the UUID of its SVM.

```
# The API:
GET /api/protocols/audit/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "ec650e97-156e-11e9-abcb-005056bbd0bf",
    "name": "vs1"
  },
  "enabled": true,
  "events": {
    "authorization_policy": false,
    "cap_staging": false,
    "cifs_logon_logoff": true,
    "file_operations": true,
    "file_share" : false,
    "security_group": false,
    "user_account": false
  },
  "log": {
    "format": "evtx",
    "rotation": {
      "size": 2048000
    },
    "retention": {
      "count": 10,
      "duration": "0s"
    }
  },
  "log_path": "/"
}
```

Updating a specific audit configuration of an SVM

The configuration is identified by the UUID of its SVM and the provided information is updated.

```
# The API:
PATCH /api/protocols/audit/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf" -H "accept: application/json" -H "Content-Type:
application/json" -d '{"enabled": false}'
```

Deleting a specific audit configuration for an SVM

The entry to be deleted is identified by the UUID of its SVM.

```
# The API:
DELETE /api/protocols/audit/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/audit/ec650e97-156e-11e9-
abcb-005056bbd0bf" -H "accept: application/json"
```

Retrieve audit configurations

GET /protocols/audit

Introduced In: 9.6

Retrieves audit configurations.

Related ONTAP commands

- `vserver audit show`

Learn more

- [DOC /protocols/audit](#)

Parameters

Name	Type	In	Required	Description
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name

Name	Type	In	Required	Description
events.file_operations	boolean	query	False	Filter by events.file_operations
events.security_group	boolean	query	False	Filter by events.security_group
events.cifs_logon_logoff	boolean	query	False	Filter by events.cifs_logon_logoff
events.authorization_policy	boolean	query	False	Filter by events.authorization_policy
events.user_account	boolean	query	False	Filter by events.user_account
events.file_share	boolean	query	False	Filter by events.file_share
events.cap_staging	boolean	query	False	Filter by events.cap_staging
log_path	string	query	False	Filter by log_path
log.retention.duration	string	query	False	Filter by log.retention.duration
log.retention.count	integer	query	False	Filter by log.retention.count
log.format	string	query	False	Filter by log.format
log.rotation.schedule.months	integer	query	False	Filter by log.rotation.schedule.months
log.rotation.schedule.days	integer	query	False	Filter by log.rotation.schedule.days

Name	Type	In	Required	Description
log.rotation.schedule.minutes	integer	query	False	Filter by log.rotation.schedule.minutes
log.rotation.schedule.hours	integer	query	False	Filter by log.rotation.schedule.hours
log.rotation.schedule.weekdays	integer	query	False	Filter by log.rotation.schedule.weekdays
log.rotation.size	integer	query	False	Filter by log.rotation.size
enabled	boolean	query	False	Filter by enabled
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[audit]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "log": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "format": "string",
        "retention": {
          "duration": "P4DT12H30M5S"
        },
        "rotation": {
          "schedule": {
            "days": [
              "integer"
            ],
            "hours": [
              "integer"
            ],
            "minutes": [
              "integer"
            ],
            "months": [
              "integer"
            ],
            "weekdays": [
              "integer"
            ]
          }
        }
      },
      "log_path": "string",
      "svm": {
        "_links": {
```

```

        "self": {
            "href": "/api/resourcelink"
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

events

Name	Type	Description
authorization_policy	boolean	Authorization policy change events
cap_staging	boolean	Central access policy staging events
cifs_logon_logoff	boolean	CIFS logon and logoff events
file_operations	boolean	File operation events
file_share	boolean	File share category events
security_group	boolean	Local security group management events
user_account	boolean	Local user account management events

_links

Name	Type	Description
self	href	

retention

Name	Type	Description
count	integer	Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with duration.
duration	string	Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with count.

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

Name	Type	Description
days	array[integer]	Specifies the day of the month schedule to rotate audit log. Leave empty for all.
hours	array[integer]	Specifies the hourly schedule to rotate audit log. Leave empty for all.
minutes	array[integer]	Specifies the minutes schedule to rotate the audit log.
months	array[integer]	Specifies the months schedule to rotate audit log. Leave empty for all.
weekdays	array[integer]	Specifies the weekdays schedule to rotate audit log. Leave empty for all.

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

Name	Type	Description
now	boolean	Manually rotates the audit logs. Optional in PATCH only. Not available in POST.
schedule	audit_schedule	Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.
size	integer	Rotates logs based on log size in bytes.

log

Name	Type	Description
_links	_links	
format	string	<p>The format in which the logs are generated by consolidation process.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • xml - Data ONTAP-specific XML log format • evtx - Microsoft Windows EVTX log format <ul style="list-style-type: none"> ◦ Default value: 1 ◦ enum: ["xml", "evtx"] ◦ Introduced in: 9.6
retention	retention	
rotation	rotation	Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

audit

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on SVMs.

Name	Type	Description
enabled	boolean	Specifies whether or not auditing is enabled on the SVM.
events	events	
log	log	
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create an audit configuration

POST /protocols/audit

Introduced In: 9.6

Creates an audit configuration.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM to which audit configuration is to be created.
- `log_path` - Path in the owning SVM namespace that is used to store audit logs.

Default property values

If not specified in POST, the following default property values are assigned:

- `enabled` - *true*
- `events.authorization_policy` - *false*
- `events.cap_staging` - *false*
- `events.file_share` - *false*
- `events.security_group` - *false*
- `events.user_account` - *false*
- `events.cifs_logon_logoff` - *true*
- `events.file_operations` - *true*
- `log.format` - *evtx*
- `log.retention.count` - *0*
- `log.retention.duration` - *PT0S*
- `log.rotation.size` - *100MB*
- `log.rotation.now` - *false*

Related ONTAP commands

- `vserver audit create`
- `vserver audit enable`

Learn more

- [DOC /protocols/audit](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
enabled	boolean	Specifies whether or not auditing is enabled on the SVM.
events	events	
log	log	

Name	Type	Description
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	svm	

Example request

```
{
  "log": {
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
num_records	integer	Number of records
records	array [audit]	

Example response

```
{
  "records": [
    {
      "log": {
        "format": "string",
        "retention": {
          "duration": "P4DT12H30M5S"
        },
        "rotation": {
          "schedule": {
            "days": [
              "integer"
            ],
            "hours": [
              "integer"
            ],
            "minutes": [
              "integer"
            ],
            "months": [
              "integer"
            ],
            "weekdays": [
              "integer"
            ]
          }
        }
      },
      "log_path": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
262196	Log_rotation_now is not an allowed operation
2621462	The specified SVM does not exist
9699330	An audit configuration already exists
9699337	Audit system internal update is in progress, audit configuration create failed
9699340	SVM UUID lookup failed
9699358	Audit configuration is absent for enabling
9699359	Audit configuration is already enabled
9699360	Final consolidation is in progress, audit enable failed
9699365	Enabling of audit configuration failed
9699370	Auditing was successfully configured, however audit configuration could not be enabled
9699384	The specified log_path does not exist
9699385	The log_path must be a directory
9699386	The log_path must be a canonical path in the SVMs namespace
9699387	The log_path cannot be empty
9699388	Rotate size must be greater than or equal to 1024 KB
9699389	The log_path must not contain a symbolic link
9699398	The log_path exceeds a maximum supported length of characters
9699399	The log_path contains an unsupported read-only (DP/LS) volume
9699400	The specified log_path is not a valid destination for SVM
9699402	The log_path contains an unsupported snaplock volume
9699403	The log_path cannot be accessed for validation
9699406	The log_path validation failed
9699409	Failed to enable multiproto.audit.evtxlog.support support capability
9699428	All nodes need to run ONTAP 8.3.0 release to audit CIFS logon-logoff events
9699429	Failed to enable multiproto.audit.cifslogonlogoff.support support capability

Error Code	Description
9699431	All nodes need to run ONTAP 8.3.0 release to audit CAP staging events
9699432	Failed to enable multiproto.audit.capstaging.support support capability

Definitions

See Definitions

events

Name	Type	Description
authorization_policy	boolean	Authorization policy change events
cap_staging	boolean	Central access policy staging events
cifs_logon_logoff	boolean	CIFS logon and logoff events
file_operations	boolean	File operation events
file_share	boolean	File share category events
security_group	boolean	Local security group management events
user_account	boolean	Local user account management events

href

Name	Type	Description
href	string	

_links

retention

Name	Type	Description
count	integer	Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with duration.
duration	string	Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with count.

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

Name	Type	Description
days	array[integer]	Specifies the day of the month schedule to rotate audit log. Leave empty for all.
hours	array[integer]	Specifies the hourly schedule to rotate audit log. Leave empty for all.
minutes	array[integer]	Specifies the minutes schedule to rotate the audit log.
months	array[integer]	Specifies the months schedule to rotate audit log. Leave empty for all.
weekdays	array[integer]	Specifies the weekdays schedule to rotate audit log. Leave empty for all.

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

Name	Type	Description
now	boolean	Manually rotates the audit logs. Optional in PATCH only. Not available in POST.
schedule	audit_schedule	Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.
size	integer	Rotates logs based on log size in bytes.

log

Name	Type	Description
format	string	<p>The format in which the logs are generated by consolidation process.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • xml - Data ONTAP-specific XML log format • evtx - Microsoft Windows EVTX log format <ul style="list-style-type: none"> ◦ Default value: 1 ◦ enum: ["xml", "evtx"] ◦ Introduced in: 9.6
retention	retention	
rotation	rotation	<p>Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.</p>

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

audit

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on SVMs.

Name	Type	Description
enabled	boolean	Specifies whether or not auditing is enabled on the SVM.
events	events	
log	log	

Name	Type	Description
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete an audit configuration

```
DELETE /protocols/audit/{svm.uuid}
```

Introduced In: 9.6

Deletes an audit configuration.

Related ONTAP commands

- `vserver audit disable`
- `vserver audit delete`

Learn more

- [DOC /protocols/audit](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
svm.uuid	string	path	True	<p>UUID of the SVM to which this object belongs.</p>

Response

```
Status: 202, Accepted
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
9699349	Auditing should be disabled before deleting the audit configuration
9699350	Audit configuration cannot be deleted, final consolidation is in progress
9699410	Failed to disable multiproto.audit.evtxlog.support support capability
9699430	Failed to disable multiproto.audit.cifslogonlogoff.support support capability
9699433	Failed to disable multiproto.audit.capstaging.support support capability

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve an audit configuration for an SVM

GET /protocols/audit/{svm.uuid}

Introduced In: 9.6

Retrieves an audit configuration for an SVM.

Related ONTAP commands

- `vserver audit show`

Learn more

- [DOC /protocols/audit](#)

Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
enabled	boolean	Specifies whether or not auditing is enabled on the SVM.
events	events	
log	log	
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	svm	

Example response

```
{
  "log": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

events

Name	Type	Description
authorization_policy	boolean	Authorization policy change events
cap_staging	boolean	Central access policy staging events
cifs_logon_logoff	boolean	CIFS logon and logoff events
file_operations	boolean	File operation events
file_share	boolean	File share category events
security_group	boolean	Local security group management events
user_account	boolean	Local user account management events

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

retention

Name	Type	Description
count	integer	Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with duration.

Name	Type	Description
duration	string	Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with count.

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

Name	Type	Description
days	array[integer]	Specifies the day of the month schedule to rotate audit log. Leave empty for all.
hours	array[integer]	Specifies the hourly schedule to rotate audit log. Leave empty for all.
minutes	array[integer]	Specifies the minutes schedule to rotate the audit log.
months	array[integer]	Specifies the months schedule to rotate audit log. Leave empty for all.
weekdays	array[integer]	Specifies the weekdays schedule to rotate audit log. Leave empty for all.

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

Name	Type	Description
now	boolean	Manually rotates the audit logs. Optional in PATCH only. Not available in POST.

Name	Type	Description
schedule	audit_schedule	Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.
size	integer	Rotates logs based on log size in bytes.

log

Name	Type	Description
_links	_links	
format	string	<p>The format in which the logs are generated by consolidation process.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • xml - Data ONTAP-specific XML log format • evtx - Microsoft Windows EVTX log format <ul style="list-style-type: none"> ◦ Default value: 1 ◦ enum: ["xml", "evtx"] ◦ Introduced in: 9.6
retention	retention	
rotation	rotation	Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

svm

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update an audit configuration for an SVM

PATCH /protocols/audit/{svm.uuid}

Introduced In: 9.6

Updates an audit configuration for an SVM.

Related ONTAP commands

- `vserver audit modify`

Learn more

- [DOC /protocols/audit](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
enabled	boolean	Specifies whether or not auditing is enabled on the SVM.
events	events	
log	log	

Name	Type	Description
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	svm	

Example request

```
{
  "log": {
    "format": "string",
    "retention": {
      "duration": "P4DT12H30M5S"
    },
    "rotation": {
      "schedule": {
        "days": [
          "integer"
        ],
        "hours": [
          "integer"
        ],
        "minutes": [
          "integer"
        ],
        "months": [
          "integer"
        ],
        "weekdays": [
          "integer"
        ]
      }
    }
  },
  "log_path": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```


Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
9699340	SVM UUID lookup failed
9699343	Audit configuration is absent for modification
9699358	Audit configuration is absent for enabling
9699359	Audit configuration is already enabled
9699360	Final consolidation is in progress, audit enable failed
9699365	Enabling of audit configuration failed
9699373	Audit configuration is absent for disabling
9699374	Audit configuration is already disabled
9699375	Disabling of audit configuration failed
9699384	The specified log_path does not exist
9699385	The log_path must be a directory
9699386	The log_path must be a canonical path in the SVMs namespace
9699387	The log_path cannot be empty
9699388	Rotate size must be greater than or equal to 1024 KB
9699389	The log_path must not contain a symbolic link
9699398	The log_path exceeds a maximum supported length of characters
9699399	The log_path contains an unsupported read-only (DP/LS) volume
9699400	The specified log_path is not a valid destination for SVM
9699402	The log_path contains an unsupported snaplock volume
9699403	The log_path cannot be accessed for validation
9699406	The log_path validation failed
9699407	Additional fields are provided

Error Code	Description
9699409	Failed to enable multiproto.audit.evtxlog.support support capability
9699410	Failed to disable multiproto.audit.evtxlog.support support capability
9699418	Audit configuration is absent for rotate
9699419	Failed to rotate audit log
9699420	Cannot rotate audit log, auditing is not enabled for this SVM
9699428	All nodes need to run ONTAP 8.3.0 release to audit CIFS logon-logoff events
9699429	Failed to enable multiproto.audit.cifslogonlogoff.support support capability
9699430	Failed to disable multiproto.audit.cifslogonlogoff.support support capability
9699431	All nodes need to run ONTAP 8.3.0 release to audit CAP staging events
9699432	Failed to enable multiproto.audit.capstaging.support support capability
9699433	Failed to disable multiproto.audit.capstaging.support support capability

Definitions

See Definitions

events

Name	Type	Description
authorization_policy	boolean	Authorization policy change events
cap_staging	boolean	Central access policy staging events
cifs_logon_logoff	boolean	CIFS logon and logoff events
file_operations	boolean	File operation events
file_share	boolean	File share category events
security_group	boolean	Local security group management events
user_account	boolean	Local user account management events

href

Name	Type	Description
href	string	

_links

retention

Name	Type	Description
count	integer	Determines how many audit log files to retain before rotating the oldest log file out. This is mutually exclusive with duration.
duration	string	Specifies an ISO-8601 format date and time to retain the audit log file. The audit log files are deleted once they reach the specified date/time. This is mutually exclusive with count.

audit_schedule

Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.

Name	Type	Description
days	array[integer]	Specifies the day of the month schedule to rotate audit log. Leave empty for all.
hours	array[integer]	Specifies the hourly schedule to rotate audit log. Leave empty for all.
minutes	array[integer]	Specifies the minutes schedule to rotate the audit log.
months	array[integer]	Specifies the months schedule to rotate audit log. Leave empty for all.
weekdays	array[integer]	Specifies the weekdays schedule to rotate audit log. Leave empty for all.

rotation

Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.

Name	Type	Description
now	boolean	Manually rotates the audit logs. Optional in PATCH only. Not available in POST.
schedule	audit_schedule	Rotates the audit logs based on a schedule by using the time-based rotation parameters in any combination. The rotation schedule is calculated by using all the time-related values.
size	integer	Rotates logs based on log size in bytes.

log

Name	Type	Description
format	string	<p>The format in which the logs are generated by consolidation process.</p> <p>Possible values are:</p> <ul style="list-style-type: none"> • xml - Data ONTAP-specific XML log format • evtx - Microsoft Windows EVTX log format <ul style="list-style-type: none"> ◦ Default value: 1 ◦ enum: ["xml", "evtx"] ◦ Introduced in: 9.6
retention	retention	
rotation	rotation	<p>Audit event log files are rotated when they reach a configured threshold log size or are on a configured schedule. When an event log file is rotated, the scheduled consolidation task first renames the active converted file to a time-stamped archive file, and then creates a new active converted event log file.</p>

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

audit

Auditing for NAS events is a security measure that enables you to track and log certain CIFS and NFS events on SVMs.

Name	Type	Description
enabled	boolean	Specifies whether or not auditing is enabled on the SVM.
events	events	
log	log	

Name	Type	Description
log_path	string	The audit log destination path where consolidated audit logs are stored.
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage home directory search paths

Protocols CIFS home-directory search-paths endpoint overview

Overview

ONTAP home directory functionality can be used to create home directories for SMB users on the CIFS server and automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user.

The home directory search path is a set of absolute paths from the root of an SVM that directs ONTAP to search for home directories. If there are multiple search paths, ONTAP tries them in the order specified until it finds a valid path. To use the CIFS home directories feature, at least one home directory search path must be added for an SVM.

Examples

Creating a home directory search path

To create a home directory search path, use the following API. Note the *return_records=true* query parameter

used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/home-directory/search-paths

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-paths?return_records=true" -H "accept: applicaion/json" -H "Content-Type: application/json" -d "{ \"path\": \"/\", \"svm\": { \"name\": \"vs1\", \"uuid\": \"a41fd873-ecf8-11e8-899d-0050568e9333\" } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "a41fd873-ecf8-11e8-899d-0050568e9333",
        "name": "vs1"
      },
      "path": "/"
    }
  ]
}
```

Retrieving the CIFS home directory search paths configuration for all SVMs in the cluster

```
# The API:
GET /protocols/cifs/home-directory/search-paths

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-paths?fields=*&return_records=true&return_timeout=15" -H "accept: application/json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "2d96f9aa-f4ce-11e8-b075-0050568e278e",
        "name": "vs1"
      },
    },
  ]
}
```

```

    "index": 1,
    "path": "/"
  },
  {
    "svm": {
      "uuid": "2d96f9aa-f4ce-11e8-b075-0050568e278e",
      "name": "vs1"
    },
    "index": 2,
    "path": "/a"
  },
  {
    "svm": {
      "uuid": "4f23449b-f4ce-11e8-b075-0050568e278e",
      "name": "vs2"
    },
    "index": 1,
    "path": "/"
  },
  {
    "svm": {
      "uuid": "4f23449b-f4ce-11e8-b075-0050568e278e",
      "name": "vs2"
    },
    "index": 2,
    "path": "/1"
  }
],
"num_records": 4
}

```

Retrieving a specific home directory searchpath configuration for an SVM

The configuration returned is identified by the UUID of its SVM and the index (position) in the list of search paths that is searched to find a home directory of a user.


```
# The API:
GET /api/protocols/home-directory/search-paths/{svm.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-paths/2d96f9aa-f4ce-11e8-b075-0050568e278e/2" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "2d96f9aa-f4ce-11e8-b075-0050568e278e",
    "name": "vs1"
  },
  "index": 2,
  "path": "/a"
}
```

Reordering a specific home directory search path in the list

An entry in the home directory search path list can be reordered to a new position by specifying the 'new_index' field. The reordered configuration is identified by the UUID of its SVM and the index.

```
# The API:
PATCH /api/protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-paths/2d96f9aa-f4ce-11e8-b075-0050568e278e/2?new_index=1" -H "accept: application/json"
```

Removing a specific home directory search path for an SVM

The entry being removed is identified by the UUID of its SVM and the index.

```
# The API:
DELETE /api/protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/home-directory/search-paths/2d96f9aa-f4ce-11e8-b075-0050568e278e/2" -H "accept: application/json"
```

Retrieve CIFS home directory search paths

GET /protocols/cifs/home-directory/search-paths

Introduced In: 9.6

Retrieves CIFS home directory search paths.

Related ONTAP commands

- `cifs server home-directory search-path show`

Learn more

- [DOC /protocols/cifs/home-directory/search-paths](#)

Parameters

Name	Type	In	Required	Description
index	integer	query	False	Filter by index
path	string	query	False	Filter by path
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cifs_search_path]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "index": 0,
      "path": "/HomeDirectory/EngDomain",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_search_path

This is a list of CIFS home directory search paths. When a CIFS client connects to a home directory share, these paths are searched in the order indicated by the position field to find the home directory of the connected CIFS client.

Name	Type	Description
index	integer	The position in the list of paths that is searched to find the home directory of the CIFS client. Not available in POST.
path	string	The file system path that is searched to find the home directory of the CIFS client.
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a home directory search path

POST /protocols/cifs/home-directory/search-paths

Introduced In: 9.6

Creates a home directory search path.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the home directory search path.
- `path` - Path in the owning SVM namespace that is used to search for home directories.

Related ONTAP commands

- `cifs server home-directory search-path add`

Learn more

- [DOC /protocols/cifs/home-directory/search-paths](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
index	integer	The position in the list of paths that is searched to find the home directory of the CIFS client. Not available in POST.
path	string	The file system path that is searched to find the home directory of the CIFS client.
svm	svm	

Example request

```
{
  "index": 0,
  "path": "/HomeDirectory/EngDomain",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[cifs_search_path]	

Example response

```
{
  "records": [
    {
      "index": 0,
      "path": "/HomeDirectory/EngDomain",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655551	Invalid home-directory search-path path
655462	The specified path is an invalid file-type

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_search_path

This is a list of CIFS home directory search paths. When a CIFS client connects to a home directory share, these paths are searched in the order indicated by the position field to find the home directory of the connected CIFS client.

Name	Type	Description
index	integer	The position in the list of paths that is searched to find the home directory of the CIFS client. Not available in POST.
path	string	The file system path that is searched to find the home directory of the CIFS client.
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a CIFS home directory search path

```
DELETE /protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}
```

Introduced In: 9.6

Deletes a CIFS home directory search path.

Related ONTAP commands

- `cifs server home-directory search-path remove`

Learn more

- [DOC /protocols/cifs/home-directory/search-paths](#)

Parameters

Name	Type	In	Required	Description
index	integer	path	True	Home directory search path index
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Response

```
Status: 200, Ok
```

Error

```
Status: Default, Error
```

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a CIFS home directory search path for an SVM

GET /protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}

Introduced In: 9.6

Retrieves a CIFS home directory search path of an SVM.

Related ONTAP commands

- `cifs server home-directory search-path show`

Learn more

- [DOC /protocols/cifs/home-directory/search-paths](#)

Parameters

Name	Type	In	Required	Description
index	integer	path	True	Home directory search path index
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
index	integer	The position in the list of paths that is searched to find the home directory of the CIFS client. Not available in POST.
path	string	The file system path that is searched to find the home directory of the CIFS client.
svm	svm	

Example response

```
{
  "index": 0,
  "path": "/HomeDirectory/EngDomain",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Re-order a CIFS home directory search path

PATCH /protocols/cifs/home-directory/search-paths/{svm.uuid}/{index}

Introduced In: 9.6

Reorders a CIFS home directory search path.

Related ONTAP commands

- `cifs server home-directory search-path reorder`

Learn more

- [DOC /protocols/cifs/home-directory/search-paths](#)

Parameters

Name	Type	In	Required	Description
index	integer	path	True	Home directory search path index
new_index	integer	query	False	New position for the home directory search path
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655463	Failed to reorder the search-path because the new-index is invalid. It cannot be '0' and it cannot go beyond the current entries

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage CIFS services

Protocols CIFS services endpoint overview

Overview

A CIFS server is necessary to provide SMB clients with access to the Storage Virtual Machine (SVM). Before you begin, the following prerequisites must be in place:

- At least one SVM LIF must exist on the SVM.
- The LIFs must be able to connect to the DNS servers configured on the SVM and to an Active Directory domain controller of the domain to which you want to join the CIFS server.
- The DNS servers must contain the service location records that are needed to locate the Active Directory domain services.
- The cluster time must be synchronized to within five minutes of the Active Directory domain controller.

Performance monitoring

Performance of the SVM can be monitored by the `metric.*` and `statistics.*` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

Information on the CIFS server

You must keep the following in mind when creating the CIFS server:

- The CIFS server name might or might not be the same as the SVM name.
- The CIFS server name can be up to 15 characters in length.
- The following characters are not allowed: @ # * () = + [] \ | ; : " , < > \ / ?
- You must use the FQDN when specifying the domain.
- The default is to add the CIFS server machine account to the Active Directory "CN=Computer" object.
- You can choose to add the CIFS server to a different organizational unit (OU) by specifying the "organizational_unit" parameter. When specifying the OU, do not specify the domain portion of the distinguished name; only specify the OU or CN portion of the distinguished name. ONTAP appends the value provided for the required "-domain" parameter onto the value provided for the "-ou" parameter to create the Active Directory distinguished name, which is used when joining the Active Directory domain.
- You can optionally choose to add a text comment of up to 48 characters about the CIFS server. If there is a space in the comment text, you must enclose the entire string in quotation marks.
- You can optionally choose to add a comma-delimited list of one or more NetBIOS aliases for the CIFS server.
- The initial administrative status of the CIFS server is "up".
- The `<i>large-mtu</i>` and `multichannel` features are enabled for the new CIFS server.
- If LDAP is configured with the `use_start_tls` and `session_security` features, the new CIFS server will also have this property set.

Examples

Creating a CIFS server

To create a CIFS server, use the following API. Note the `return_records=true` query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/services

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/services?return_records=true" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"ad_domain\": { \"fqdn\": \"CIFS-2008R2-AD.GDL.ENGLAB.NETAPP.COM\", \"organizational_unit\": \"CN=Computers\", \"password\": \"cifs*123\", \"user\": \"administrator\" }, \"comment\": \"This CIFS Server Belongs to CS Department\", \"default_unix_user\": \"string\", \"enabled\": true, \"name\": \"CIFS-DOC\", \"netbios\": { \"aliases\": [ \"ALIAS_1\", \"ALIAS_2\", \"ALIAS_3\" ], \"enabled\": false, \"wins_servers\": [ \"10.224.65.20\", \"10.224.65.21\" ] }, \"security\": { \"kdc_encryption\": false, \"encrypt_dc_connection\": false,
```

```

\"restrict_anonymous\": \"no_enumeration\", \"smb_encryption\": false,
\"smb_signing\": false }, \"svm\": { \"name\": \"vs1\", \"uuid\":
\"ef087155-f9e2-11e8-ac52-0050568ea248\" }}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "9f5ab4cb-f703-11e8-91cc-0050568eca13",
        "name": "vs1"
      },
      "name": "CIFS-DOC",
      "ad_domain": {
        "fqdn": "CIFS-2008R2-AD.GDL.ENGLAB.NETAPP.COM",
        "user": "administrator",
        "password": "cifs*123",
        "organizational_unit": "CN=Computers"
      },
      "enabled": true,
      "comment": "This CIFS Server Belongs to CS Department",
      "security": {
        "restrict_anonymous": "no_enumeration",
        "smb_signing": false,
        "smb_encryption": false,
        "encrypt_dc_connection": false,
        "kdc_encryption": false
      },
      "netbios": {
        "aliases": [
          "ALIAS_1",
          "ALIAS_2",
          "ALIAS_3"
        ],
        "wins_servers": [
          "10.224.65.20",
          "10.224.65.21"
        ],
        "enabled": false
      },
      "default_unix_user": "string"
    }
  ],
  "job": {
    "uuid": "f232b6da-00a4-11e9-a8c1-0050568eca13",

```

```

    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f232b6da-00a4-11e9-a8c1-0050568eca13"
      }
    }
  }
}

```

Retrieving the full CIFS server configuration for all SVMs in the cluster

```

# The API:
GET /api/protocols/cifs/services

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/cifs/services?fields=*&return_records=true&return_timeou
t=15" -H "accept: application/json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "9f5ab4cb-f703-11e8-91cc-0050568eca13",
        "name": "vs1"
      },
      "name": "CIFS-DOC",
      "ad_domain": {
        "fqdn": "CIFS-2008R2-AD.GDL.ENGLAB.NETAPP.COM",
        "organizational_unit": "CN=Computers"
      },
      "enabled": true,
      "comment": "This CIFS Server Belongs to CS Department",
      "security": {
        "restrict_anonymous": "no_enumeration",
        "smb_signing": false,
        "smb_encryption": false,
        "encrypt_dc_connection": false,
        "kdc_encryption": false,
        "lm_compatibility_level": "lm_ntlm_ntlmv2_krb"
      },
      "netbios": {

```

```
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ],
    "enabled": false
  },
  "default_unix_user": "string"
}
],
"num_records": 1
}
```

Retrieving CIFS server configuration details for a specific SVM

```
# The API:
GET /api/protocols/cifs/services/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/services/9f5ab4cb-f703-11e8-91cc-0050568eca13" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "9f5ab4cb-f703-11e8-91cc-0050568eca13",
    "name": "vs1"
  },
  "name": "CIFS-DOC",
  "ad_domain": {
    "fqdn": "CIFS-2008R2-AD.GDL.ENGLAB.NETAPP.COM",
    "organizational_unit": "CN=Computers"
  },
  "enabled": true,
  "comment": "This CIFS Server Belongs to CS Department",
  "security": {
    "restrict_anonymous": "no_enumeration",
    "smb_signing": false,
    "smb_encryption": false,
    "encrypt_dc_connection": false,
    "kdc_encryption": false,
    "lm_compatibility_level": "lm_ntlm_ntlmv2_krb"
  },
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ],
    "enabled": false
  },
  "default_unix_user": "string"
}
```

Updating CIFS server properties for the specified SVM

```
# The API:
PATCH /api/protocols/cifs/services/{svm.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/services/9f5ab4cb-f703-11e8-91cc-0050568eca13" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"comment\": \"CIFS SERVER MODIFICATION\" }"
```

Removing a CIFS server for a specific SVM

To delete a CIFS server, use the following API. This will delete the CIFS server along with other CIFS configurations such as CIFS share, share ACLs, homedir search-path, and so on.

```
# The API:
DELETE /api/protocols/cifs/services/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/services/9f5ab4cb-f703-11e8-91cc-0050568eca13" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"ad_domain\": { \"password\": \"cifs*123\", \"user\": \"administrator\" } }"
```

Retrieve CIFS servers

GET /protocols/cifs/services

Introduced In: 9.6

Retrieves CIFS servers.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `statistics.*`
- `metric.*`

Related ONTAP commands

- `vserver cifs server show`
- `vserver cifs server options show`
- `vserver cifs server security show`

Learn more

- [DOC /protocols/cifs/services](#)

Parameters

Name	Type	In	Required	Description
metric.throughput.write	integer	query	False	Filter by metric.throughput.write • Introduced in: 9.7
metric.throughput.read	integer	query	False	Filter by metric.throughput.read • Introduced in: 9.7
metric.throughput.total	integer	query	False	Filter by metric.throughput.total • Introduced in: 9.7
metric.duration	string	query	False	Filter by metric.duration • Introduced in: 9.7
metric.timestamp	string	query	False	Filter by metric.timestamp • Introduced in: 9.7

Name	Type	In	Required	Description
metric.status	string	query	False	Filter by metric.status • Introduced in: 9.7
metric.iops.total	integer	query	False	Filter by metric.iops.total • Introduced in: 9.7
metric.iops.read	integer	query	False	Filter by metric.iops.read • Introduced in: 9.7
metric.iops.other	integer	query	False	Filter by metric.iops.other • Introduced in: 9.7
metric.iops.write	integer	query	False	Filter by metric.iops.write • Introduced in: 9.7
metric.latency.total	integer	query	False	Filter by metric.latency.total • Introduced in: 9.7
metric.latency.read	integer	query	False	Filter by metric.latency.read • Introduced in: 9.7
metric.latency.other	integer	query	False	Filter by metric.latency.other • Introduced in: 9.7

Name	Type	In	Required	Description
metric.latency.write	integer	query	False	Filter by metric.latency.write • Introduced in: 9.7
comment	string	query	False	Filter by comment
name	string	query	False	Filter by name
default_unix_user	string	query	False	Filter by default_unix_user
ad_domain.user	string	query	False	Filter by ad_domain.user
ad_domain.fqdn	string	query	False	Filter by ad_domain.fqdn
ad_domain.organizational_unit	string	query	False	Filter by ad_domain.organizational_unit
enabled	boolean	query	False	Filter by enabled
security.encrypt_dc_connection	boolean	query	False	Filter by security.encrypt_dc_connection • Introduced in: 9.8
security.lm_compatibility_level	string	query	False	Filter by security.lm_compatibility_level • Introduced in: 9.8
security.smb_encryption	boolean	query	False	Filter by security.smb_encryption
security.kdc_encryption	boolean	query	False	Filter by security.kdc_encryption

Name	Type	In	Required	Description
security.smb_signing	boolean	query	False	Filter by security.smb_signing
security.restrict_anonymous	string	query	False	Filter by security.restrict_anonymous
netbios.enabled	boolean	query	False	Filter by netbios.enabled
netbios.wins_servers	string	query	False	Filter by netbios.wins_servers
netbios.aliaes	string	query	False	Filter by netbios.aliaes
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
statistics.latency_raw.total	integer	query	False	Filter by statistics.latency_raw.total • Introduced in: 9.7
statistics.latency_raw.read	integer	query	False	Filter by statistics.latency_raw.read • Introduced in: 9.7
statistics.latency_raw.other	integer	query	False	Filter by statistics.latency_raw.other • Introduced in: 9.7

Name	Type	In	Required	Description
statistics.latency_raw.write	integer	query	False	Filter by statistics.latency_raw.write • Introduced in: 9.7
statistics.timestamp	string	query	False	Filter by statistics.timestamp • Introduced in: 9.7
statistics.iops_raw.total	integer	query	False	Filter by statistics.iops_raw.total • Introduced in: 9.7
statistics.iops_raw.read	integer	query	False	Filter by statistics.iops_raw.read • Introduced in: 9.7
statistics.iops_raw.other	integer	query	False	Filter by statistics.iops_raw.other • Introduced in: 9.7
statistics.iops_raw.write	integer	query	False	Filter by statistics.iops_raw.write • Introduced in: 9.7
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write • Introduced in: 9.7

Name	Type	In	Required	Description
statistics.throughput_raw.read	integer	query	False	Filter by statistics.throughput_raw.read • Introduced in: 9.7
statistics.throughput_raw.total	integer	query	False	Filter by statistics.throughput_raw.total • Introduced in: 9.7
statistics.status	string	query	False	Filter by statistics.status • Introduced in: 9.7
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cifs_service]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ad_domain": {
        "fqdn": "example.com",
        "organizational_unit": "string",
        "user": "string"
      },
      "comment": "This CIFS Server Belongs to CS Department",
      "default_unix_user": "string",
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "latency": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "status": "ok",
        "throughput": {
          "read": "200",
          "total": "1000",
```

```

        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"name": "CIFS1",
"netbios": {
    "aliases": [
        "ALIAS_1",
        "ALIAS_2",
        "ALIAS_3"
    ],
    "wins_servers": [
        "10.224.65.20",
        "10.224.65.21"
    ]
},
"security": {
    "lm_compatibility_level": "string",
    "restrict_anonymous": "string"
},
"statistics": {
    "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",

```



```
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

ad_domain

Name	Type	Description
fqdn	string	The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only.
user	string	The user account used to add this CIFS server to the Active Directory. POST and DELETE only.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Peformance metric for write I/O operations.

metric

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

cifs_netbios

Name	Type	Description
aliases	array[string]	
enabled	boolean	Specifies whether NetBios name service (NBNS) is enabled for the CIFS. If this service is enabled, the CIFS server will start sending the broadcast for name registration.
wins_servers	array[string]	

cifs_service_security

Name	Type	Description
encrypt_dc_connection	boolean	Specifies whether encryption is required for domain controller connections.

Name	Type	Description
kdc_encryption	boolean	<p>Specifies whether AES-128 and AES-256 encryption is enabled for all Kerberos-based communication with the Active Directory KDC.</p> <p>To take advantage of the strongest security with Kerberos-based communication, AES-256 and AES-128 encryption can be enabled on the CIFS server.</p> <p>Kerberos-related communication for CIFS is used during CIFS server creation on the SVM, as well as during the SMB session setup phase.</p> <p>The CIFS server supports the following encryption types for Kerberos communication:</p> <ul style="list-style-type: none"> • RC4-HMAC • DES • AES <p>When the CIFS server is created, the domain controller creates a computer machine account in Active Directory. After a newly created machine account authenticates, the KDC and the CIFS server negotiates encryption types. At this time, the KDC becomes aware of the encryption capabilities of the particular machine account and uses those capabilities in subsequent communication with the CIFS server.</p> <p>In addition to negotiating encryption types during CIFS server creation, the encryption types are renegotiated when a machine account password is reset.</p>

Name	Type	Description
lm_compatibility_level	string	<p>It is CIFS server minimum security level, also known as the LMCompatibilityLevel. The minimum security level is the minimum level of the security tokens that the CIFS server accepts from SMB clients. The available values are:</p> <ul style="list-style-type: none"> • lm_ntlm_ntlmv2_krb Accepts LM, NTLM, NTLMv2 and Kerberos • ntlm_ntlmv2_krb Accepts NTLM, NTLMv2 and Kerberos • ntlmv2_krb Accepts NTLMv2 and Kerberos • krb Accepts Kerberos only
restrict_anonymous	string	<p>Specifies what level of access an anonymous user is granted. An anonymous user (also known as a "null user") can list or enumerate certain types of system information from Windows hosts on the network, including user names and details, account policies, and share names. Access for the anonymous user can be controlled by specifying one of three access restriction settings. The available values are:</p> <ul style="list-style-type: none"> • no_restriction - No access restriction for an anonymous user. • no_enumeration - Enumeration is restricted for an anonymous user. • no_access - All access is restricted for an anonymous user.
smb_encryption	boolean	Specifies whether encryption is required for incoming CIFS traffic.

Name	Type	Description
smb_signing	boolean	Specifies whether signing is required for incoming CIFS traffic. SMB signing helps to ensure that network traffic between the CIFS server and the client is not compromised.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_service

Name	Type	Description
_links	_links	

Name	Type	Description
ad_domain	ad_domain	
comment	string	A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks.
default_unix_user	string	Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails.
enabled	boolean	Specifies if the CIFS service is administratively enabled.
metric	metric	
name	string	The name of the CIFS server.
netbios	cifs_netbios	
security	cifs_service_security	
statistics	statistics	
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Create a CIFS server

POST /protocols/cifs/services

Introduced In: 9.6

Creates a CIFS server. Each SVM can have one CIFS server.

Important notes

- The CIFS server name might or might not be the same as the SVM name.
- The CIFS server name can contain up to 15 characters.
- The CIFS server name does not support the following characters: @ # * () = + [] \ | ; : " , < > / ?

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the CIFS server.
- `name` - Name of the CIFS server.
- `ad_domain.fqdn` - Fully qualified domain name of the Windows Active Directory to which this CIFS server belongs.
- `ad_domain.user` - User account with the access to add the CIFS server to the Active Directory.
- `ad_domain.password` - Account password used to add this CIFS server to the Active Directory.

Recommended optional properties

- `comment` - Add a text comment of up to 48 characters about the CIFS server.
- `netbios.aliases` - Add a comma-delimited list of one or more NetBIOS aliases for the CIFS server.
- `netbios.wins_servers` - Add a list of Windows Internet Name Server (WINS) addresses that manage and map the NetBIOS name of the CIFS server to their network IP addresses. The IP addresses must be IPv4 addresses.

Default property values

If not specified in POST, the following default property values are assigned:

- `ad_domain.organizational_unit` - *CN=Computers*
- `enabled` - *true*
- `restrict_anonymous` - *no_enumeration*
- `smb_signing` - *false*
- `smb_encryption` - *false*

- `encrypt_dc_connection` - *false*
- `kdc_encryption` - *false*
- `default_unix_user` - *pcuser*
- `netbios_enabled` - *false* However, if either "netbios.wins-server" or "netbios.aliases" is set during POST and if `netbios_enabled` is not specified then `netbios_enabled` is set to true.

Related ONTAP commands

- `vserver cifs server create`
- `vserver cifs server options modify`
- `vserver cifs security modify`
- `vserver cifs server add-netbios-aliases`

Learn more

- [DOC /protocols/cifs/services](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
ad_domain	ad_domain	

Name	Type	Description
comment	string	A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks.
default_unix_user	string	Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails.
enabled	boolean	Specifies if the CIFS service is administratively enabled.
metric	metric	
name	string	The name of the CIFS server.
netbios	cifs_netbios	
security	cifs_service_security	
statistics	statistics	
svm	svm	

Example request

```
{
  "ad_domain": {
    "fqdn": "example.com",
    "organizational_unit": "string",
    "password": "string",
    "user": "string"
  },
  "comment": "This CIFS Server Belongs to CS Department",
  "default_unix_user": "string",
  "metric": {
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "name": "CIFS1",
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ]
  },
  "security": {
    "restrict_anonymous": "string"
  },
}
```



```

"statistics": {
  "iops_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

{
  "job": {
    "uuid": "string"
  }
}

```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

ad_domain

Name	Type	Description
fqdn	string	The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only.
password	string	The account password used to add this CIFS server to the Active Directory. This is not audited.
user	string	The user account used to add this CIFS server to the Active Directory. POST and DELETE only.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.

Name	Type	Description
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

cifs_netbios

Name	Type	Description
aliases	array[string]	
enabled	boolean	Specifies whether NetBios name service (NBNS) is enabled for the CIFS. If this service is enabled, the CIFS server will start sending the broadcast for name registration.
wins_servers	array[string]	

cifs_service_security

Name	Type	Description
encrypt_dc_connection	boolean	Specifies whether encryption is required for domain controller connections.

Name	Type	Description
kdc_encryption	boolean	<p>Specifies whether AES-128 and AES-256 encryption is enabled for all Kerberos-based communication with the Active Directory KDC.</p> <p>To take advantage of the strongest security with Kerberos-based communication, AES-256 and AES-128 encryption can be enabled on the CIFS server.</p> <p>Kerberos-related communication for CIFS is used during CIFS server creation on the SVM, as well as during the SMB session setup phase.</p> <p>The CIFS server supports the following encryption types for Kerberos communication:</p> <ul style="list-style-type: none"> • RC4-HMAC • DES • AES <p>When the CIFS server is created, the domain controller creates a computer machine account in Active Directory. After a newly created machine account authenticates, the KDC and the CIFS server negotiates encryption types. At this time, the KDC becomes aware of the encryption capabilities of the particular machine account and uses those capabilities in subsequent communication with the CIFS server.</p> <p>In addition to negotiating encryption types during CIFS server creation, the encryption types are renegotiated when a machine account password is reset.</p>

Name	Type	Description
restrict_anonymous	string	<p>Specifies what level of access an anonymous user is granted. An anonymous user (also known as a "null user") can list or enumerate certain types of system information from Windows hosts on the network, including user names and details, account policies, and share names. Access for the anonymous user can be controlled by specifying one of three access restriction settings.</p> <p>The available values are:</p> <ul style="list-style-type: none"> • no_restriction - No access restriction for an anonymous user. • no_enumeration - Enumeration is restricted for an anonymous user. • no_access - All access is restricted for an anonymous user.
smb_encryption	boolean	Specifies whether encryption is required for incoming CIFS traffic.
smb_signing	boolean	Specifies whether signing is required for incoming CIFS traffic. SMB signing helps to ensure that network traffic between the CIFS server and the client is not compromised.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_service

Name	Type	Description
ad_domain	ad_domain	
comment	string	A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks.
default_unix_user	string	Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails.
enabled	boolean	Specifies if the CIFS service is administratively enabled.
metric	metric	
name	string	The name of the CIFS server.
netbios	cifs_netbios	
security	cifs_service_security	
statistics	statistics	
svm	svm	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a CIFS server and related configurations

```
DELETE /protocols/cifs/services/{svm.uuid}
```

Introduced In: 9.6

Deletes a CIFS server and related CIFS configurations.

Related ONTAP commands

- `vserver cifs server delete`
- `vserver cifs remove-netbios-aliases`

Learn more

- [DOC /protocols/cifs/services](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
ad_domain	ad_domain	

Example request

```
{
  "ad_domain": {
    "fqdn": "example.com",
    "organizational_unit": "string",
    "password": "string",
    "user": "string"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

ad_domain

Name	Type	Description
fqdn	string	The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only.
password	string	The account password used to add this CIFS server to the Active Directory. This is not audited.
user	string	The user account used to add this CIFS server to the Active Directory. POST and DELETE only.

cifs_service_delete

Name	Type	Description
ad_domain	ad_domain	

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a CIFS server

GET /protocols/cifs/services/{svm.uuid}

Introduced In: 9.6

Retrieves a CIFS server.

Related ONTAP commands

- `vserver cifs server show`
- `vserver cifs server options show`
- `vserver cifs server security show`

Learn more

- [DOC /protocols/cifs/services](#)

Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
ad_domain	ad_domain	
comment	string	A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks.
default_unix_user	string	Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails.
enabled	boolean	Specifies if the CIFS service is administratively enabled.
metric	metric	
name	string	The name of the CIFS server.
netbios	cifs_netbios	
security	cifs_service_security	
statistics	statistics	
svm	svm	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_domain": {
    "fqdn": "example.com",
    "organizational_unit": "string",
    "user": "string"
  },
  "comment": "This CIFS Server Belongs to CS Department",
  "default_unix_user": "string",
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"name": "CIFS1",
"netbios": {
  "aliases": [
    "ALIAS_1",
    "ALIAS_2",
    "ALIAS_3"
  ]
}
```

```

    ],
    "wins_servers": [
        "10.224.65.20",
        "10.224.65.21"
    ]
},
"security": {
    "lm_compatibility_level": "string",
    "restrict_anonymous": "string"
},
"statistics": {
    "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

ad_domain

Name	Type	Description
fqdn	string	The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only.
user	string	The user account used to add this CIFS server to the Active Directory. POST and DELETE only.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.

Name	Type	Description
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Name	Type	Description
<code>_links</code>	_links	

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

cifs_netbios

Name	Type	Description
aliases	array[string]	
enabled	boolean	Specifies whether NetBios name service (NBNS) is enabled for the CIFS. If this service is enabled, the CIFS server will start sending the broadcast for name registration.
wins_servers	array[string]	

cifs_service_security

Name	Type	Description
encrypt_dc_connection	boolean	Specifies whether encryption is required for domain controller connections.

Name	Type	Description
kdc_encryption	boolean	<p>Specifies whether AES-128 and AES-256 encryption is enabled for all Kerberos-based communication with the Active Directory KDC.</p> <p>To take advantage of the strongest security with Kerberos-based communication, AES-256 and AES-128 encryption can be enabled on the CIFS server.</p> <p>Kerberos-related communication for CIFS is used during CIFS server creation on the SVM, as well as during the SMB session setup phase.</p> <p>The CIFS server supports the following encryption types for Kerberos communication:</p> <ul style="list-style-type: none"> • RC4-HMAC • DES • AES <p>When the CIFS server is created, the domain controller creates a computer machine account in Active Directory. After a newly created machine account authenticates, the KDC and the CIFS server negotiates encryption types. At this time, the KDC becomes aware of the encryption capabilities of the particular machine account and uses those capabilities in subsequent communication with the CIFS server.</p> <p>In addition to negotiating encryption types during CIFS server creation, the encryption types are renegotiated when a machine account password is reset.</p>

Name	Type	Description
lm_compatibility_level	string	<p>It is CIFS server minimum security level, also known as the LMCompatibilityLevel. The minimum security level is the minimum level of the security tokens that the CIFS server accepts from SMB clients. The available values are:</p> <ul style="list-style-type: none"> • lm_ntlm_ntlmv2_krb Accepts LM, NTLM, NTLMv2 and Kerberos • ntlm_ntlmv2_krb Accepts NTLM, NTLMv2 and Kerberos • ntlmv2_krb Accepts NTLMv2 and Kerberos • krb Accepts Kerberos only
restrict_anonymous	string	<p>Specifies what level of access an anonymous user is granted. An anonymous user (also known as a "null user") can list or enumerate certain types of system information from Windows hosts on the network, including user names and details, account policies, and share names. Access for the anonymous user can be controlled by specifying one of three access restriction settings. The available values are:</p> <ul style="list-style-type: none"> • no_restriction - No access restriction for an anonymous user. • no_enumeration - Enumeration is restricted for an anonymous user. • no_access - All access is restricted for an anonymous user.
smb_encryption	boolean	Specifies whether encryption is required for incoming CIFS traffic.

Name	Type	Description
smb_signing	boolean	Specifies whether signing is required for incoming CIFS traffic. SMB signing helps to ensure that network traffic between the CIFS server and the client is not compromised.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update CIFS mandatory and optional parameters

PATCH /protocols/cifs/services/{svm.uuid}

Introduced In: 9.6

Updates both the mandatory and optional parameters of the CIFS configuration. Ensure the CIFS server is administratively disabled when renaming the CIFS server or modifying the *ad_domain* properties.

Related ONTAP commands

- `vserver cifs server modify`
- `vserver cifs server options modify`
- `vserver cifs security modify`
- `vserver cifs server add-netbios-aliases`
- `vserver cifs server remove-netbios-aliases`

Learn more

- [DOC /protocols/cifs/services](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
ad_domain	ad_domain	
comment	string	A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks.

Name	Type	Description
default_unix_user	string	Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails.
enabled	boolean	Specifies if the CIFS service is administratively enabled.
metric	metric	
name	string	The name of the CIFS server.
netbios	cifs_netbios	
security	cifs_service_security	
statistics	statistics	
svm	svm	

Example request

```
{
  "ad_domain": {
    "fqdn": "example.com",
    "organizational_unit": "string",
    "password": "string"
  },
  "comment": "This CIFS Server Belongs to CS Department",
  "default_unix_user": "string",
  "metric": {
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "name": "CIFS1",
  "netbios": {
    "aliases": [
      "ALIAS_1",
      "ALIAS_2",
      "ALIAS_3"
    ],
    "wins_servers": [
      "10.224.65.20",
      "10.224.65.21"
    ]
  },
  "security": {
    "restrict_anonymous": "string"
  },
  "statistics": {
```

```
"iops_raw": {
  "read": "200",
  "total": "1000",
  "write": "100"
},
"latency_raw": {
  "read": "200",
  "total": "1000",
  "write": "100"
},
"status": "ok",
"throughput_raw": {
  "read": "200",
  "total": "1000",
  "write": "100"
},
"timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

ad_domain

Name	Type	Description
fqdn	string	The fully qualified domain name of the Windows Active Directory to which this CIFS server belongs. A CIFS server appears as a member of Windows server object in the Active Directory store. POST and PATCH only.
organizational_unit	string	Specifies the organizational unit within the Active Directory domain to associate with the CIFS server. POST and PATCH only.
password	string	The account password used to add this CIFS server to the Active Directory. This is not audited.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

cifs_netbios

Name	Type	Description
aliases	array[string]	
enabled	boolean	Specifies whether NetBios name service (NBNS) is enabled for the CIFS. If this service is enabled, the CIFS server will start sending the broadcast for name registration.
wins_servers	array[string]	

cifs_service_security

Name	Type	Description
encrypt_dc_connection	boolean	Specifies whether encryption is required for domain controller connections.

Name	Type	Description
kdc_encryption	boolean	<p>Specifies whether AES-128 and AES-256 encryption is enabled for all Kerberos-based communication with the Active Directory KDC.</p> <p>To take advantage of the strongest security with Kerberos-based communication, AES-256 and AES-128 encryption can be enabled on the CIFS server.</p> <p>Kerberos-related communication for CIFS is used during CIFS server creation on the SVM, as well as during the SMB session setup phase.</p> <p>The CIFS server supports the following encryption types for Kerberos communication:</p> <ul style="list-style-type: none"> • RC4-HMAC • DES • AES <p>When the CIFS server is created, the domain controller creates a computer machine account in Active Directory. After a newly created machine account authenticates, the KDC and the CIFS server negotiates encryption types. At this time, the KDC becomes aware of the encryption capabilities of the particular machine account and uses those capabilities in subsequent communication with the CIFS server.</p> <p>In addition to negotiating encryption types during CIFS server creation, the encryption types are renegotiated when a machine account password is reset.</p>

Name	Type	Description
restrict_anonymous	string	<p>Specifies what level of access an anonymous user is granted. An anonymous user (also known as a "null user") can list or enumerate certain types of system information from Windows hosts on the network, including user names and details, account policies, and share names. Access for the anonymous user can be controlled by specifying one of three access restriction settings.</p> <p>The available values are:</p> <ul style="list-style-type: none"> • no_restriction - No access restriction for an anonymous user. • no_enumeration - Enumeration is restricted for an anonymous user. • no_access - All access is restricted for an anonymous user.
smb_encryption	boolean	Specifies whether encryption is required for incoming CIFS traffic.
smb_signing	boolean	Specifies whether signing is required for incoming CIFS traffic. SMB signing helps to ensure that network traffic between the CIFS server and the client is not compromised.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_service

Name	Type	Description
ad_domain	ad_domain	
comment	string	A descriptive text comment for the CIFS server. SMB clients can see the CIFS server comment when browsing servers on the network. If there is a space in the comment, you must enclose the entire string in quotation marks.
default_unix_user	string	Specifies the UNIX user to which any authenticated CIFS user is mapped to, if the normal user mapping rules fails.
enabled	boolean	Specifies if the CIFS service is administratively enabled.
metric	metric	
name	string	The name of the CIFS server.
netbios	cifs_netbios	
security	cifs_service_security	
statistics	statistics	
svm	svm	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve CIFS protocol historical performance metrics for an SVM

GET /protocols/cifs/services/{svm.uuid}/metrics

Introduced In: 9.7

Retrieves historical performance metrics for the CIFS protocol of an SVM.

Parameters

Name	Type	In	Required	Description
throughput.total	integer	query	False	Filter by throughput.total
throughput.read	integer	query	False	Filter by throughput.read
throughput.other	integer	query	False	Filter by throughput.other
throughput.write	integer	query	False	Filter by throughput.write
timestamp	string	query	False	Filter by timestamp
duration	string	query	False	Filter by duration

Name	Type	In	Required	Description
status	string	query	False	Filter by status
iops.total	integer	query	False	Filter by iops.total
iops.read	integer	query	False	Filter by iops.read
iops.other	integer	query	False	Filter by iops.other
iops.write	integer	query	False	Filter by iops.write
latency.total	integer	query	False	Filter by latency.total
latency.read	integer	query	False	Filter by latency.read
latency.other	integer	query	False	Filter by latency.other
latency.write	integer	query	False	Filter by latency.write
svm.uuid	string	path	True	Unique identifier of the SVM.

Name	Type	In	Required	Description
interval	string	query	False	<p>The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y. The period for each time range is as follows:</p> <ul style="list-style-type: none"> • 1h: Metrics over the most recent hour sampled over 15 seconds. • 1d: Metrics over the most recent day sampled over 5 minutes. • 1w: Metrics over the most recent week sampled over 30 minutes. • 1m: Metrics over the most recent month sampled over 2 hours. • 1y: Metrics over the most recent year sampled over a day. • Default value: 1 • enum: ["1h", "1d", "1w", "1m", "1y"]

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

records

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View CIFS active sessions

Protocols CIFS sessions endpoint overview

Overview

ONTAP CIFS sessions `show` functionality is used to provide a list of currently established CIFS sessions.

The following lists the fields retrieved using the CIFS sessions GET API:

node.name: Node name hosting this record; basically the node hosting the "server_ip".
node.uuid: Node UUID hosting this record; basically the node hosting the "server_ip".
svm.name: Svm name to which the "server_ip" belongs to.
svm.uuid: Svm uuid to which the "server_ip" belongs to.
server_ip: All clients connected to this interface are displayed in rows.
client_ip: IP address of the client connected to the interface.
identifier: Unique identifier used to represent each SMB session.
connection_id: Unique identifier used to represent each SMB Connection.
authentication: Type of authentication supported by the server when a client accesses a SMB share.
user: Username for the Windows client.
mapped_unix_user: Mapped UNIX user name.
open_shares: Number of shares opened by the client on a specific SVM.
open_files: Number of files opened by the client on a specific SVM.
open_others: Number of other files opened by the client on a specific SVM.
connected_duration: Time elapsed since the first request was sent by the client for this SMB session.
idle_duration: Time elapsed since the last request was sent by the client for this SMB session.
protocol: SMB protocol dialects over which the client accesses the SMB share.
availability: Level of continuous availability of protection provided to the files from the SMB share
smb_signing: Specifies whether SMB signing is enabled.
smb_encryption: Specifies the SMB session encryption status.
large_mtu: Specifies whether the SMB session's large MTU is enabled.

connection_count: Counter used to track requests that are sent to the volumes to the node.

volumes.name: Name of the active volumes the client is accessing.

volumes.uuid: UUID of the active volumes the client is accessing.

Example

Retrieves established sessions information

```
# The API:
GET /protocols/cifs/sessions

# The call:
curl -X GET "https://<cluster-mgmt-
ip>/api/protocols/cifs/sessions?return_timeout=15&return_records=true" -H
"accept: application/json"

# The response:
{
  "records": [
    {
      "node": {
        "uuid": "85d46998-4e5d-11ea-afb1-0050568ec4e4",
        "name": "bkalyan-vsim1"
      },
      "svm": {
        "uuid": "fc824aa8-4e60-11ea-afb1-0050568ec4e4",
        "name": "vs1"
      },
      "identifier": 625718873227788300,
      "connection_id": 91842,
      "lif_address": "10.140.70.197",
      "address": "10.74.7.182",
      "auth_mechanism": "ntlmv2",
      "windows_user": "NBCIFSQA2\\administrator",
      "unix_user": "root",
      "shares": 1,
      "files": 2,
      "other": 0,
      "connected_time": "PT16H54M47S",
      "idle_time": "PT3S",
      "protocol_version": "smb3_1",
      "continuously_available": "no",
      "is_session_signed": false,
      "smb_encryption_status": "unencrypted",
      "connection_count": 1,
      "is_large_mtu_enabled": true,
      "vol_names": [
```



```

        "vol12",
        "origin",
        "fg",
        "vol1"
    ],
    "vol_uuids": [
        "954d697f-4e62-11ea-afb1-0050568ec4e4",
        "0f909e06-4e6e-11ea-afb1-0050568ec4e4",
        "08c27b7d-4e61-11ea-afb1-0050568ec4e4",
        "02f0a240-4e61-11ea-afb1-0050568ec4e4"
    ]
},
{
    "node": {
        "uuid": "85d46998-4e5d-11ea-afb1-0050568ec4e4",
        "name": "bkalyan-vsim1"
    },
    "svm": {
        "uuid": "fc824aa8-4e60-11ea-afb1-0050568ec4e4",
        "name": "vs1"
    },
    "identifier": 625718873227788500,
    "connection_id": 92080,
    "lif_address": "10.140.70.197",
    "address": "10.140.133.97",
    "auth_mechanism": "ntlmv2",
    "windows_user": "NBCIFSQA2\\administrator",
    "unix_user": "root",
    "shares": 1,
    "files": 1,
    "other": 0,
    "connected_time": "PT1M19S",
    "idle_time": "PT1M1S",
    "protocol_version": "smb3",
    "continuously_available": "no",
    "is_session_signed": false,
    "smb_encryption_status": "unencrypted",
    "connection_count": 1,
    "is_large_mtu_enabled": true,
    "vol_names": [
        "origin",
        "fg",
        "vol1"
    ],
    "vol_uuids": [
        "0f909e06-4e6e-11ea-afb1-0050568ec4e4",

```

```

    "08c27b7d-4e61-11ea-afb1-0050568ec4e4",
    "02f0a240-4e61-11ea-afb1-0050568ec4e4"
  ]
},
"num_records": 2
}

```

Retrieve CIFS session information for all SVMs

GET /protocols/cifs/sessions

Introduced In: 9.8

Retrieves the CIFS sessions information for all SVMs.

Related ONTAP commands

- `vserver cifs session show -active-volumes`

Parameters

Name	Type	In	Required	Description
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
large_mtu	boolean	query	False	Filter by large_mtu
authentication	string	query	False	Filter by authentication
user	string	query	False	Filter by user
open_other	integer	query	False	Filter by open_other
client_ip	string	query	False	Filter by client_ip
smb_encryption	string	query	False	Filter by smb_encryption
continuous_availability	string	query	False	Filter by continuous_availability
connected_duration	string	query	False	Filter by connected_duration

Name	Type	In	Required	Description
mapped_unix_user	string	query	False	Filter by mapped_unix_user
idle_duration	string	query	False	Filter by idle_duration
volumes.uuid	string	query	False	Filter by volumes.uuid
volumes.name	string	query	False	Filter by volumes.name
open_shares	integer	query	False	Filter by open_shares
server_ip	string	query	False	Filter by server_ip
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
connection_id	integer	query	False	Filter by connection_id
smb_signing	boolean	query	False	Filter by smb_signing
open_files	integer	query	False	Filter by open_files
identifier	integer	query	False	Filter by identifier
protocol	string	query	False	Filter by protocol
connection_count	integer	query	False	Filter by connection_count
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cifs_session]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication": "ntlmv2",
      "client_ip": "10.74.7.182",
      "connected_duration": "P4DT84H30M5S",
      "connection_count": "0",
      "connection_id": "22802",
      "continuous_availability": "no",
      "identifier": "4622663542519103507",
      "idle_duration": "P4DT84H30M5S",
      "large_mtu": 1,
      "mapped_unix_user": "root",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
      },
      "open_files": 0,
      "open_other": 0,
      "open_shares": 0,
      "protocol": "smb3.1",
      "server_ip": "10.140.78.248",
      "smb_encryption": "unencrypted",
      "smb_signing": "",
      "svm": {
        "_links": {
```

```

        "self": {
            "href": "/api/resourcelink"
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "user": "NBCIFSQA2\\administrator",
    "volumes": [
        {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
    ]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volumes

Name	Type	Description
_links	_links	
name	string	The name of the volume.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7

cifs_session

Name	Type	Description
_links	_links	
authentication	string	<p>SMB authentication over which the client accesses the share. The following values are supported:</p> <ul style="list-style-type: none"> none - No authentication ntlmv1 - Ntlm version 1 mechanism ntlmv2 - Ntlm version 2 mechanism kerberos - Kerberos authentication anonymous - Anonymous mechanism
client_ip	string	Specifies IP address of the client.
connected_duration	string	Specifies an ISO-8601 format of date and time used to retrieve the connected time duration in hours, minutes, and seconds format.
connection_count	integer	A counter used to track requests that are sent to the volumes to the node.
connection_id	integer	A unique identifier used to represent each SMB connection.

Name	Type	Description
continuous_availability	string	<p>The level of continuous availability protection provided to the SMB sessions and/or files.</p> <ul style="list-style-type: none"> • no - Open file is not continuously available. For sessions, it contains one or more open files but none of them are continuously available. • yes - open file is continuously available. For sessions, it contains one or more open files and all of them are continuously available. • partial - Sessions only. Contains at least one continuously available open file with other files open but not continuously available.
identifier	integer	A unique identifier used to represent each SMB session.
idle_duration	string	Specifies an ISO-8601 format of date and time used to retrieve the idle time duration in hours, minutes, and seconds format.
large_mtu	boolean	Specifies whether or not a large MTU is enabled for an SMB session.
mapped_unix_user	string	Indicated that a mapped UNIX user has logged in.
node	node	
open_files	integer	Number of files the SMB session has opened.
open_other	integer	Number of other files the SMB session has opened.
open_shares	integer	Number of shares the SMB session has opened.

Name	Type	Description
protocol	string	<p>The SMB protocol version over which the client accesses the volumes. The following values are supported:</p> <ul style="list-style-type: none"> • smb1 - SMB version 1 • smb2 - SMB version 2 • smb2.1 - SMB version 2 minor version 1 • smb3 - SMB version 3 • smb3.1 - SMB version 3 minor version 1
server_ip	string	Specifies the IP address of the SVM.
smb_encryption	string	<p>Indicates an SMB encryption state. The following values are supported:</p> <ul style="list-style-type: none"> • unencrypted - SMB session is not encrypted • encrypted - SMB session is fully encrypted. SVM level encryption is enabled and encryption occurs for the entire session. • partially_encrypted - SMB session is partially encrypted. Share level encryption is enabled and encryption is initiated when the tree-connect occurs.
smb_signing	boolean	Specifies whether or not SMB signing is enabled.
svm	svm	
user	string	Indicates that a Windows user has logged in.
volumes	array[volumes]	A group of volumes, the client is accessing.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve an SMB session information for an SMB connection in a SVM node

GET /protocols/cifs/sessions/{node.uuid}/{svm.uuid}/{identifier}/{connection_id}

Introduced In: 9.8

Retrieves SMB session information for a specific SMB connection of a SVM in a node.

Learn more

- [DOC /protocols/cifs/sessions](#)

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID.
identifier	integer	path	True	Unique identifier for the SMB session.
connection_id	integer	path	True	Unique identifier for the SMB connection.
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
authentication	string	<p>SMB authentication over which the client accesses the share. The following values are supported:</p> <ul style="list-style-type: none"> • none - No authentication • ntlmv1 - Ntlm version 1 mechanism • ntlmv2 - Ntlm version 2 mechanism • kerberos - Kerberos authentication • anonymous - Anonymous mechanism
client_ip	string	Specifies IP address of the client.
connected_duration	string	Specifies an ISO-8601 format of date and time used to retrieve the connected time duration in hours, minutes, and seconds format.
connection_count	integer	A counter used to track requests that are sent to the volumes to the node.
connection_id	integer	A unique identifier used to represent each SMB connection.

Name	Type	Description
continuous_availability	string	<p>The level of continuous availability protection provided to the SMB sessions and/or files.</p> <ul style="list-style-type: none"> • no - Open file is not continuously available. For sessions, it contains one or more open files but none of them are continuously available. • yes - open file is continuously available. For sessions, it contains one or more open files and all of them are continuously available. • partial - Sessions only. Contains at least one continuously available open file with other files open but not continuously available.
identifier	integer	A unique identifier used to represent each SMB session.
idle_duration	string	Specifies an ISO-8601 format of date and time used to retrieve the idle time duration in hours, minutes, and seconds format.
large_mtu	boolean	Specifies whether or not a large MTU is enabled for an SMB session.
mapped_unix_user	string	Indicated that a mapped UNIX user has logged in.
node	node	
open_files	integer	Number of files the SMB session has opened.
open_other	integer	Number of other files the SMB session has opened.
open_shares	integer	Number of shares the SMB session has opened.

Name	Type	Description
protocol	string	<p>The SMB protocol version over which the client accesses the volumes. The following values are supported:</p> <ul style="list-style-type: none"> • smb1 - SMB version 1 • smb2 - SMB version 2 • smb2.1 - SMB version 2 minor version 1 • smb3 - SMB version 3 • smb3.1 - SMB version 3 minor version 1
server_ip	string	Specifies the IP address of the SVM.
smb_encryption	string	<p>Indicates an SMB encryption state. The following values are supported:</p> <ul style="list-style-type: none"> • unencrypted - SMB session is not encrypted • encrypted - SMB session is fully encrypted. SVM level encryption is enabled and encryption occurs for the entire session. • partially_encrypted - SMB session is partially encrypted. Share level encryption is enabled and encryption is initiated when the tree-connect occurs.
smb_signing	boolean	Specifies whether or not SMB signing is enabled.
svm	svm	
user	string	Indicates that a Windows user has logged in.
volumes	array[volumes]	A group of volumes, the client is accessing.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication": "ntlmv2",
  "client_ip": "10.74.7.182",
  "connected_duration": "P4DT84H30M5S",
  "connection_count": "0",
  "connection_id": "22802",
  "continuous_availability": "no",
  "identifier": "4622663542519103507",
  "idle_duration": "P4DT84H30M5S",
  "large_mtu": 1,
  "mapped_unix_user": "root",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "open_files": 0,
  "open_other": 0,
  "open_shares": 0,
  "protocol": "smb3.1",
  "server_ip": "10.140.78.248",
  "smb_encryption": "unencrypted",
  "smb_signing": "",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "user": "NBCIFSQA2\\administrator",
  "volumes": [
    {
```



```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volumes

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage CIFS shares

Protocols CIFS shares endpoint overview

Overview

Before any users or applications can access data on the CIFS server over SMB, a CIFS share must be created with sufficient share permissions. CIFS share is a named access point in a volume which is tied to the CIFS server on the SVM. Before creating a CIFS share make sure that the path is valid within the scope of the SVM and that it is reachable.

Permissions can be assigned to this newly created share by specifying the 'acls' field. When a CIFS share is created, ONTAP creates a default ACL for this share with 'Full-Control' permissions for an 'Everyone' user.

Examples

Creating a CIFS share

To create a CIFS share for a CIFS server, use the following API. Note the *return_records=true* query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/shares

# The call:
curl -X POST "https://<mgmt-
ip>/api/protocols/cifs/shares?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{
\"access_based_enumeration\": false, \"acls\": [ { \"permission\":
\"no_access\", \"type\": \"unix_user\", \"user_or_group\": \"root\" } ],
\"change_notify\": true, \"comment\": \"HR Department Share\",
\"encryption\": false, \"home_directory\": false, \"name\": \"TEST\",
\"oplocks\": true, \"path\": \"/\", \"svm\": { \"name\": \"vs1\",
\"uuid\": \"000c5cd2-ebdf-11e8-a96e-0050568ea3cb\" }, \"unix_symlink\":
\"local\"}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
        "name": "vs1"
      },
      "name": "TEST",
      "path": "/",
      "comment": "HR Department Share",
      "home_directory": false,
      "oplocks": true,
      "access_based_enumeration": false,
      "change_notify": true,
      "encryption": false,
      "unix_symlink": "local",
      "acls": [
        {
          "user_or_group": "root",
          "type": "unix_user",
          "permission": "no_access",
          "winsid_unixId": "0"
        }
      ]
    }
  ]
}
```

Retrieving CIFS Shares for all SVMs in the cluster

```
# The API:
GET /api/protocols/cifs/shares

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/cifs/shares?fields=*&return_records=true&return_timeout=
15" -H "accept application/hal+json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/000c5cd2-ebdf-11e8-a96e-0050568ea3cb"
          }
        }
      },
      "name": "admin$",
      "path": "/",
      "home_directory": false,
      "oplocks": false,
      "access_based_enumeration": false,
      "change_notify": false,
      "encryption": false,
      "volume": {
        "name": "vol1",
        "uuid": "4e06f1bc-1ddc-42e2-abb2-f221c6a2ab2a"
      },
      "_links": {
        "self": {
          "href": "/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-
0050568ea3cb/admin%24"
        }
      }
    },
    {
      "svm": {
        "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
```

```

    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/000c5cd2-ebdf-11e8-a96e-0050568ea3cb"
      }
    }
  },
  "name": "c$",
  "path": "/",
  "home_directory": false,
  "oplocks": true,
  "access_based_enumeration": false,
  "change_notify": true,
  "encryption": false,
  "unix_symlink": "local",
  "acls": [
    {
      "user_or_group": "BUILTIN\\Administrators",
      "type": "windows",
      "permission": "full_control"
    }
  ],
  "volume": {
    "name": "vol1",
    "uuid": "4e06f1bc-1ddc-42e2-abb2-f221c6a2ab2a"
  },
  "_links": {
    "self": {
      "href": "/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/c%24"
    }
  }
},
{
  "svm": {
    "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/000c5cd2-ebdf-11e8-a96e-0050568ea3cb"
      }
    }
  },
  "name": "ipc$",
  "path": "/",
  "home_directory": false,

```

```

    "oplocks": false,
    "access_based_enumeration": false,
    "change_notify": false,
    "encryption": false,
    "volume": {
      "name": "vol1",
      "uuid": "4e06f1bc-1ddc-42e2-abb2-f221c6a2ab2a"
    },
    "_links": {
      "self": {
        "href": "/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/ipc%24"
      }
    }
  },
  {
    "svm": {
      "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/000c5cd2-ebdf-11e8-a96e-0050568ea3cb"
        }
      }
    },
    "name": "TEST",
    "path": "/",
    "comment": "HR Department Share",
    "home_directory": false,
    "oplocks": true,
    "access_based_enumeration": false,
    "change_notify": true,
    "encryption": false,
    "unix_symlink": "local",
    "acls": [
      {
        "user_or_group": "Everyone",
        "type": "windows",
        "permission": "full_control"
      },
      {
        "user_or_group": "root",
        "type": "unix_user",
        "permission": "no_access"
      }
    ]
  },

```

```
    "volume": {
      "name": "vol1",
      "uuid": "4e06f1bc-1ddc-42e2-abb2-f221c6a2ab2a"
    },
    "_links": {
      "self": {
        "href": "/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/TEST"
      }
    }
  ],
  "num_records": 4,
  "_links": {
    "self": {
      "href":
"/api/protocols/cifs/shares?fields=*&return_records=true&return_timeout=15"
    }
  }
}
```

Retrieving all CIFS Shares for all SVMs in the cluster for which the acls are configured for a "root" user

```

# The API:
GET /api/protocols/cifs/shares

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/cifs/shares?acls.user_or_group=root&fields=*&return_reco
rds=true&return_timeout=15" -H "accept: application/json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
        "name": "vs1"
      },
      "name": "TEST",
      "path": "/",
      "comment": "HR Department Share",
      "home_directory": false,
      "oplocks": true,
      "access_based_enumeration": false,
      "change_notify": true,
      "encryption": false,
      "unix_symlink": "local",
      "acls": [
        {
          "user_or_group": "Everyone",
          "type": "windows",
          "permission": "full_control"
        },
        {
          "user_or_group": "root",
          "type": "unix_user",
          "permission": "no_access"
        }
      ],
      "volume": {
        "name": "vol1",
        "uuid": "4e06f1bc-1ddc-42e2-abb2-f221c6a2ab2a"
      }
    }
  ],
  "num_records": 1
}

```

Retrieving a specific CIFS share configuration for an SVM

The configuration being returned is identified by the UUID of its SVM and the name of the share.

```
# The API:
GET /api/protocols/cifs/shares/{svm.uuid}/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/TEST" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
    "name": "vs1"
  },
  "name": "TEST",
  "path": "/",
  "comment": "HR Department Share",
  "home_directory": false,
  "oplocks": true,
  "access_based_enumeration": false,
  "change_notify": true,
  "encryption": false,
  "unix_symlink": "local",
  "acls": [
    {
      "user_or_group": "Everyone",
      "type": "windows",
      "permission": "full_control"
    },
    {
      "user_or_group": "root",
      "type": "unix_user",
      "permission": "no_access"
    }
  ],
  "volume": {
    "name": "vol1",
    "uuid": "4e06f1bc-1ddc-42e2-abb2-f221c6a2ab2a"
  }
}
```

Updating a specific CIFS share for an SVM

The CIFS share being modified is identified by the UUID of its SVM and the CIFS share name. The CIFS share ACLs cannot be modified with this API.

```
# The API:
PATCH /api/protocols/cifs/shares/{svm.uuid}/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/TEST" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"access_based_enumeration\": true, \"change_notify\": true, \"comment\": \"HR Department Share\", \"encryption\": false, \"oplocks\": true, \"path\": \"/\", \"unix_symlink\": \"widelink\"}"
```

Removing a specific CIFS share for an SVM

The CIFS share being removed is identified by the UUID of its SVM and the CIFS share name.

```
# The API:
DELETE /api/protocols/cifs/shares/{svm.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/test" -H "accept: application/json"
```

Retrieve CIFS shares

GET /protocols/cifs/shares

Introduced In: 9.6

Retrieves CIFS shares.

Related ONTAP commands

- `vserver cifs share show`
- `vserver cifs share properties show`

Learn more

- [DOC /protocols/cifs/shares](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
path	string	query	False	Filter by path
change_notify	boolean	query	False	Filter by change_notify
home_directory	boolean	query	False	Filter by home_directory
access_based_enumeration	boolean	query	False	Filter by access_based_enumeration
name	string	query	False	Filter by name
comment	string	query	False	Filter by comment
acls.permission	string	query	False	Filter by acls.permission
acls.user_or_group	string	query	False	Filter by acls.user_or_group
acls.type	string	query	False	Filter by acls.type
unix_symlink	string	query	False	Filter by unix_symlink
oplocks	boolean	query	False	Filter by oplocks
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
encryption	boolean	query	False	Filter by encryption
fields	array[string]	query	False	Specify the fields to return.

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cifs_share]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "acls": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "permission": "string",
          "type": "string",
          "user_or_group": "ENGDOMAIN\\ad_user"
        }
      ],
      "comment": "HR Department Share",
      "name": "HR_SHARE",
      "path": "/volume_1/eng_vol/",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "unix_symlink": "string",
      "volume": {
        "_links": {
          "self": {
```

```

        "href": "/api/resourcelink"
      },
    ],
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

cifs_share_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	_links	
permission	string	<p>Specifies the access rights that a user or group has on the defined CIFS Share.</p> <p>The following values are allowed:</p> <ul style="list-style-type: none">• no_access - User does not have CIFS share access• read - User has only read access• change - User has change access• full_control - User has full_control access

Name	Type	Description
type	string	Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed: <ul style="list-style-type: none"> • windows - Windows user or group • unix_user - UNIX user • unix_group - UNIX group
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB, a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server

on the SVM.

When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

Name	Type	Description
_links	_links	
access_based_enumeration	boolean	If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents the display of folders or other shared resources that the user does not have access to.
acls	array[cifs_share_acl]	
change_notify	boolean	Specifies whether CIFS clients can request for change notifications for directories on this share.
comment	string	Specify the CIFS share descriptions.
encryption	boolean	Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share.

Name	Type	Description
home_directory	boolean	<p>Specifies whether or not the share is a home directory share, where the share and path names are dynamic.</p> <p>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user.</p> <p>The ONTAP CIFS home directory feature enable us to configure a share that maps to different directories based on the user that connects to it. Instead of creating a separate shares for each user, a single share with a home directory parameters can be created.</p> <p>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting %w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.</p> <ul style="list-style-type: none"> • Default value: • Introduced in: 9.6 • readCreate: 1
name	string	<p>Specifies the name of the CIFS share that you want to create. If this is a home directory share then the share name includes the pattern as %w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically.</p>

Name	Type	Description
oplocks	boolean	Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations.
path	string	<p>The fully-qualified pathname in the owning SVM namespace that is shared through this share. If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination. ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.</p> <ul style="list-style-type: none"> • example: /volume_1/eng_vol/ • maxLength: 256 • minLength: 1 • Introduced in: 9.6
svm	svm	
unix_symlink	string	<p>Controls the access of UNIX symbolic links to CIFS clients. The supported values are:</p> <ul style="list-style-type: none"> • local - Enables only local symbolic links which is within the same CIFS share. • widelink - Enables both local symlinks and widelinks. • disable - Disables local symlinks and widelinks.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a CIFS share

POST /protocols/cifs/shares

Introduced In: 9.6

Creates a CIFS share.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the CIFS share.
- `name` - Name of the CIFS share.
- `path` - Path in the owning SVM namespace that is shared through this share.

Recommended optional properties

- `comment` - Optionally choose to add a text comment of up to 256 characters about the CIFS share.
- `acls` - Optionally choose to add share permissions that users and groups have on the CIFS share.

Default property values

If not specified in POST, the following default property values are assigned:

- `home_directory` - *false*
- `oplocks` - *true*
- `access_based_enumeration` - *false*
- `change_notify` - *true*

- `encryption` - *false*
- `unix_symlink` - *local*

Related ONTAP commands

- `vserver cifs share create`
- `vserver cifs share properties add`
- `vserver cifs share access-control create`

Learn more

- [DOC /protocols/cifs/shares](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
access_based_enumeration	boolean	If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents the display of folders or other shared resources that the user does not have access to.
acls	array[cifs_share_acl]	
change_notify	boolean	Specifies whether CIFS clients can request for change notifications for directories on this share.
comment	string	Specify the CIFS share descriptions.
encryption	boolean	Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share.

Name	Type	Description
home_directory	boolean	<p>Specifies whether or not the share is a home directory share, where the share and path names are dynamic.</p> <p>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user. The ONTAP CIFS home directory feature enable us to configure a share that maps to different directories based on the user that connects to it. Instead of creating a separate shares for each user, a single share with a home directory parameters can be created.</p> <p>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting %w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.</p> <ul style="list-style-type: none"> • Default value: 1 • Introduced in: 9.6 • readCreate: 1
name	string	<p>Specifies the name of the CIFS share that you want to create. If this is a home directory share then the share name includes the pattern as %w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically.</p>
oplocks	boolean	<p>Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations.</p>

Name	Type	Description
path	string	<p>The fully-qualified pathname in the owning SVM namespace that is shared through this share.</p> <p>If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination.</p> <p>ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.</p> <ul style="list-style-type: none"> • example: /volume_1/eng_vol/ • maxLength: 256 • minLength: 1 • Introduced in: 9.6
svm	svm	
unix_symlink	string	<p>Controls the access of UNIX symbolic links to CIFS clients. The supported values are:</p> <ul style="list-style-type: none"> • local - Enables only local symbolic links which is within the same CIFS share. • widelink - Enables both local symlinks and widelinks. • disable - Disables local symlinks and widelinks.
volume	volume	

Example request

```
{
  "acls": [
    {
      "permission": "string",
      "type": "string",
      "user_or_group": "ENGDOMAIN\\ad_user"
    }
  ],
  "comment": "HR Department Share",
  "name": "HR_SHARE",
  "path": "/volume_1/eng_vol/",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_symlink": "string",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655628	CIFS Share Creation with property 'SMB_ENCRYPTION' failed because the CIFS server does not support SMB3.0
655551	CIFS Share Creation failed because the specified path does not exist
655577	The CIFS share name cannot be more than 80 characters long

Error Code	Description
655399	Failed to create CIFS share. The CIFS server does not exist for specified SVM
656422	Failed to create the home directory share because the directory shares must specify a path relative to one or more home directory search paths
656423	Failed to create CIFS share. The Shares must define an absolute share path
656424	Failed to create CIFS the administrator share 'c\$' because you are not permitted to created any admin shares
655625	Failed to create CIFS share. The Shares path is not a valid file-type for CIFS share
656426	CIFS Share Creation failed because the share name is invalid

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cifs_share_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
permission	string	Specifies the access rights that a user or group has on the defined CIFS Share. The following values are allowed: <ul style="list-style-type: none">• no_access - User does not have CIFS share access• read - User has only read access• change - User has change access• full_control - User has full_control access
type	string	Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed: <ul style="list-style-type: none">• windows - Windows user or group• unix_user - UNIX user• unix_group - UNIX group
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

svm

Name	Type	Description
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB, a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.

When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

Name	Type	Description
access_based_enumeration	boolean	If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents the display of folders or other shared resources that the user does not have access to.
acls	array[cifs_share_acl]	
change_notify	boolean	Specifies whether CIFS clients can request for change notifications for directories on this share.
comment	string	Specify the CIFS share descriptions.

Name	Type	Description
encryption	boolean	Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share.
home_directory	boolean	<p>Specifies whether or not the share is a home directory share, where the share and path names are dynamic.</p> <p>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user.</p> <p>The ONTAP CIFS home directory feature enable us to configure a share that maps to different directories based on the user that connects to it. Instead of creating a separate shares for each user, a single share with a home directory parameters can be created.</p> <p>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting %w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.</p> <ul style="list-style-type: none"> • Default value: 1 • Introduced in: 9.6 • readCreate: 1

Name	Type	Description
name	string	Specifies the name of the CIFS share that you want to create. If this is a home directory share then the share name includes the pattern as %w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically.
oplocks	boolean	Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations.
path	string	<p>The fully-qualified pathname in the owning SVM namespace that is shared through this share. If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination. ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.</p> <ul style="list-style-type: none"> • example: /volume_1/eng_vol/ • maxLength: 256 • minLength: 1 • Introduced in: 9.6
svm	svm	

Name	Type	Description
unix_symlink	string	Controls the access of UNIX symbolic links to CIFS clients. The supported values are: <ul style="list-style-type: none"> • local - Enables only local symbolic links which is within the same CIFS share. • widelink - Enables both local symlinks and widelinks. • disable - Disables local symlinks and widelinks.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a CIFS share

DELETE /protocols/cifs/shares/{svm.uuid}/{name}

Introduced In: 9.6

Deletes a CIFS share.

Related ONTAP commands

- `vserver cifs share delete`

Learn more

- [DOC /protocols/cifs/shares](#)

Parameters

Name	Type	In	Required	Description
name	string	path	True	Share Name
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655393	Standard admin shares cannot be removed

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a CIFS share

GET /protocols/cifs/shares/{svm.uuid}/{name}

Introduced In: 9.6

Retrieves a CIFS share.

Related ONTAP commands

- `vserver cifs share show`
- `vserver cifs share properties show`

Learn more

- [DOC /protocols/cifs/shares](#)

Parameters

Name	Type	In	Required	Description
name	string	path	True	Share Name
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
access_based_enumeration	boolean	If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents the display of folders or other shared resources that the user does not have access to.
acls	array[cifs_share_acl]	
change_notify	boolean	Specifies whether CIFS clients can request for change notifications for directories on this share.
comment	string	Specify the CIFS share descriptions.

Name	Type	Description
encryption	boolean	Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share.
home_directory	boolean	<p>Specifies whether or not the share is a home directory share, where the share and path names are dynamic.</p> <p>ONTAP home directory functionality automatically offer each user a dynamic share to their home directory without creating an individual SMB share for each user. The ONTAP CIFS home directory feature enable us to configure a share that maps to different directories based on the user that connects to it. Instead of creating a separate shares for each user, a single share with a home directory parameters can be created.</p> <p>In a home directory share, ONTAP dynamically generates the share-name and share-path by substituting %w, %u, and %d variables with the corresponding Windows user name, UNIX user name, and domain name, respectively.</p> <ul style="list-style-type: none"> • Default value: 1 • Introduced in: 9.6 • readCreate: 1
name	string	<p>Specifies the name of the CIFS share that you want to create. If this is a home directory share then the share name includes the pattern as %w (Windows user name), %u (UNIX user name) and %d (Windows domain name) variables in any combination with this parameter to generate shares dynamically.</p>

Name	Type	Description
oplocks	boolean	Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations.
path	string	<p>The fully-qualified pathname in the owning SVM namespace that is shared through this share.</p> <p>If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination.</p> <p>ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.</p> <ul style="list-style-type: none"> • example: /volume_1/eng_vol/ • maxLength: 256 • minLength: 1 • Introduced in: 9.6
svm	svm	
unix_symlink	string	<p>Controls the access of UNIX symbolic links to CIFS clients.</p> <p>The supported values are:</p> <ul style="list-style-type: none"> • local - Enables only local symbolic links which is within the same CIFS share. • widelink - Enables both local symlinks and widelinks. • disable - Disables local symlinks and widelinks.
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "acls": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "permission": "string",
      "type": "string",
      "user_or_group": "ENGDOMAIN\\ad_user"
    }
  ],
  "comment": "HR Department Share",
  "name": "HR_SHARE",
  "path": "/volume_1/eng_vol/",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_symlink": "string",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cifs_share_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	_links	
permission	string	<p>Specifies the access rights that a user or group has on the defined CIFS Share.</p> <p>The following values are allowed:</p> <ul style="list-style-type: none">• no_access - User does not have CIFS share access• read - User has only read access• change - User has change access• full_control - User has full_control access
type	string	<p>Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed:</p> <ul style="list-style-type: none">• windows - Windows user or group• unix_user - UNIX user• unix_group - UNIX group
user_or_group	string	<p>Specifies the user or group name to add to the access control list of a CIFS share.</p>

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a CIFS share

PATCH /protocols/cifs/shares/{svm.uuid}/{name}

Introduced In: 9.6

Updates a CIFS share.

Related ONTAP commands

- `vserver cifs share modify`
- `vserver cifs share properties add`
- `vserver cifs share properties remove`

Learn more

- [DOC /protocols/cifs/shares](#)

Parameters

Name	Type	In	Required	Description
name	string	path	True	Share Name
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
access_based_enumeration	boolean	If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents the display of folders or other shared resources that the user does not have access to.
change_notify	boolean	Specifies whether CIFS clients can request for change notifications for directories on this share.
comment	string	Specify the CIFS share descriptions.

Name	Type	Description
encryption	boolean	Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share.
oplocks	boolean	Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations.
path	string	<p>The fully-qualified pathname in the owning SVM namespace that is shared through this share. If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination. ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.</p> <ul style="list-style-type: none"> • example: /volume_1/eng_vol/ • maxLength: 256 • minLength: 1 • Introduced in: 9.6
svm	svm	
unix_symlink	string	<p>Controls the access of UNIX symbolic links to CIFS clients. The supported values are:</p> <ul style="list-style-type: none"> • local - Enables only local symbolic links which is within the same CIFS share. • widelink - Enables both local symlinks and widelinks. • disable - Disables local symlinks and widelinks.
volume	volume	

Example request

```
{
  "comment": "HR Department Share",
  "home_directory": null,
  "path": "/volume_1/eng_vol/",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_symlink": "string",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655628	'SMB_ENCRYPTION' property cannot be set on CIFS share because the CIFS server does not support SMB3.0
655551	CIFS Share modification failed because the specified path does not exist
655620	Cannot set symlink properties for admin shares
656420	Cannot modify the standard share ipc\$
656421	Cannot modify the standard share admin\$
656422	Failed to modify the home directory share because the directory shares must specify a path relative to one or more home directory search paths
656423	Failed to modify CIFS share. The Shares must define an absolute share path

Error Code	Description
656425	Failed to modify the CIFS share because the path for an administrative share cannot be modified

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cifs_share_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
permission	string	<p>Specifies the access rights that a user or group has on the defined CIFS Share.</p> <p>The following values are allowed:</p> <ul style="list-style-type: none">• no_access - User does not have CIFS share access• read - User has only read access• change - User has change access• full_control - User has full_control access

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

cifs_share

CIFS share is a named access point in a volume. Before users and applications can access data on the CIFS server over SMB, a CIFS share must be created with sufficient share permission. CIFS shares are tied to the CIFS server on the SVM.

When a CIFS share is created, ONTAP creates a default ACL for the share with Full Control permissions for Everyone.

Name	Type	Description
access_based_enumeration	boolean	If enabled, all folders inside this share are visible to a user based on that individual user access right; prevents the display of folders or other shared resources that the user does not have access to.
change_notify	boolean	Specifies whether CIFS clients can request for change notifications for directories on this share.
comment	string	Specify the CIFS share descriptions.
encryption	boolean	Specifies that SMB encryption must be used when accessing this share. Clients that do not support encryption are not able to access this share.
oplocks	boolean	Specify whether opportunistic locks are enabled on this share. "Oplocks" allow clients to lock files and cache content locally, which can increase performance for file operations.

Name	Type	Description
path	string	<p>The fully-qualified pathname in the owning SVM namespace that is shared through this share. If this is a home directory share then the path should be dynamic by specifying the pattern %w (Windows user name), %u (UNIX user name), or %d (domain name) variables in any combination. ONTAP generates the path dynamically for the connected user and this path is appended to each search path to find the full Home Directory path.</p> <ul style="list-style-type: none"> • example: /volume_1/eng_vol/ • maxLength: 256 • minLength: 1 • Introduced in: 9.6
svm	svm	
unix_symlink	string	<p>Controls the access of UNIX symbolic links to CIFS clients. The supported values are:</p> <ul style="list-style-type: none"> • local - Enables only local symbolic links which is within the same CIFS share. • widelink - Enables both local symlinks and widelinks. • disable - Disables local symlinks and widelinks.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage share-level ACL

Protocols CIFS shares svm.uuid share acls endpoint overview

Overview

Access to files and folders can be secured over a network by configuring share access control lists (ACLs) on CIFS shares. Share-level ACLs can be configured by using either Windows users and groups or UNIX users and groups. A share-level ACL consists of a list of access control entries (ACEs). Each ACE contains a user or group name and a set of permissions that determines user or group access to the share, regardless of the security style of the volume or qtree containing the share.

When an SMB user tries to access a share, ONTAP checks the share-level ACL to determine whether access should be granted. A share-level ACL only restricts access to files in the share; it never grants more access than the file level ACLs.

Examples

Creating a CIFS share ACL

To create a share ACL for a CIFS share, use the following API. Note the *return_records=true* query parameter used to obtain the newly created entry in the response.


```
# The API:
POST /api/protocols/cifs/shares{svm.uuid}/{share}/acls

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/sh1/acls?return_records=true" -H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"permission\": \"no_access\", \"type\": \"windows\", \"user_or_group\": \"root\"}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "vs1"
      },
      "user_or_group": "root",
      "type": "windows",
      "permission": "no_access"
    }
  ]
}
```

Retrieving all CIFS shares ACLs for a specific CIFS share for a specific SVM in the cluster

```
# The API:
GET /api/protocols/cifs/shares/{svm.uuid}/{share}/acls

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/sh1/acls?fields=*&return_records=true&return_timeout=15" -H "accept: application/json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
        "name": "vs1"
      },
      "share": "sh1",
      "user_or_group": "Everyone",
      "type": "windows",
      "permission": "full_control"
    },
    {
      "svm": {
        "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
        "name": "vs1"
      },
      "share": "sh1",
      "user_or_group": "root",
      "type": "windows",
      "permission": "no_access"
    }
  ],
  "num_records": 2
}
```

Retrieving a CIFS share ACLs for a user or a group of type Windows or type UNIX on a CIFS share for a specific SVM

```
# The API:
GET
/api/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/sh1/acls/everyone/windows" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
    "name": "vs1"
  },
  "share": "sh1",
  "user_or_group": "everyone",
  "type": "windows",
  "permission": "full_control"
}
```

Updating a CIFS share ACLs of a user or group on a CIFS share for a specific SVM

The CIFS share ACL being modified is identified by the UUID of its SVM, the CIFS share name, user or group name and the type of the user or group.

```
# The API:
PATCH
/api/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/sh1/acls/everyone/windows" -H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"permission\": \"no_access\""
}
```

Removing a CIFS share ACLs of a user or group on a CIFS Share for a specific SVM

The CIFS share ACL being removed is identified by the UUID of its SVM, the CIFS share name, user or group name and the type of the user or group.

```
# The API:
DELETE
/api/protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/shares/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/sh1/acls/everyone/windows" -H "accept: application/json"
```

Retrieve a share-level ACL on a CIFS share

GET /protocols/cifs/shares/{svm.uuid}/{share}/acls

Introduced In: 9.6

Retrieves the share-level ACL on a CIFS share.

Related ONTAP commands

- `vserver cifs share access-control show`

Learn more

- [DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls](#)

Parameters

Name	Type	In	Required	Description
share	string	path	True	CIFS Share Name
permission	string	query	False	Filter by permission
user_or_group	string	query	False	Filter by user_or_group
type	string	query	False	Filter by type
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
fields	array[string]	query	False	Specify the fields to return.

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cifs_share_acl]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "permission": "string",
      "type": "string",
      "user_or_group": "ENGDOMAIN\\ad_user"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

cifs_share_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
_links	_links	
permission	string	<p>Specifies the access rights that a user or group has on the defined CIFS Share.</p> <p>The following values are allowed:</p> <ul style="list-style-type: none">• no_access - User does not have CIFS share access• read - User has only read access• change - User has change access• full_control - User has full_control access

Name	Type	Description
type	string	Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed: <ul style="list-style-type: none"> • windows - Windows user or group • unix_user - UNIX user • unix_group - UNIX group
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a share-level ACL on a CIFS share

POST /protocols/cifs/shares/{svm.uuid}/{share}/acls

Introduced In: 9.6

Creates a share-level ACL on a CIFS share.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the share acl.
- `share` - Existing CIFS share in which to create the share acl.
- `user_or_group` - Existing user or group name for which the acl is added on the CIFS share.
- `permission` - Access rights that a user or group has on the defined CIFS share.

Default property values

- `type` - *windows*

Related ONTAP commands

- `vserver cifs share access-control create`

Learn more

- [DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls](#)

Parameters

Name	Type	In	Required	Description
share	string	path	True	CIFS Share Name
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
permission	string	<p>Specifies the access rights that a user or group has on the defined CIFS Share.</p> <p>The following values are allowed:</p> <ul style="list-style-type: none"> • no_access - User does not have CIFS share access • read - User has only read access • change - User has change access • full_control - User has full_control access
type	string	<p>Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed:</p> <ul style="list-style-type: none"> • windows - Windows user or group • unix_user - UNIX user • unix_group - UNIX group
user_or_group	string	<p>Specifies the user or group name to add to the access control list of a CIFS share.</p>

Example request

```
{
  "permission": "string",
  "type": "string",
  "user_or_group": "ENGDOMAIN\\ad_user"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655470	Failed to create share ACL because the share does not exist
655446	Failed to create share ACL because the specified Windows user/group does not exist
4849678	Failed to create share ACL because the specified UNIX user/group does not exist

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cifs_share_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
permission	string	Specifies the access rights that a user or group has on the defined CIFS Share. The following values are allowed: <ul style="list-style-type: none">• no_access - User does not have CIFS share access• read - User has only read access• change - User has change access• full_control - User has full_control access
type	string	Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed: <ul style="list-style-type: none">• windows - Windows user or group• unix_user - UNIX user• unix_group - UNIX group
user_or_group	string	Specifies the user or group name to add to the access control list of a CIFS share.

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a share-level ACL on a CIFS share

```
DELETE /protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}
```

Introduced In: 9.6

Deletes a share-level ACL on a CIFS share.

Related ONTAP commands

- `vserver cifs share access-control delete`

Learn more

- [DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls](#)

Parameters

Name	Type	In	Required	Description
share	string	path	True	Share name
user_or_group	string	path	True	User or group name
type	string	path	True	User or group type
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a share-level ACL on a CIFS share for a user or group

GET /protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}

Introduced In: 9.6

Retrieves the share-level ACL on CIFS share for a specified user or group.

Related ONTAP commands

- `vserver cifs share access-control show`

Learn more

- [DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls](#)

Parameters

Name	Type	In	Required	Description
share	string	path	True	Share name
user_or_group	string	path	True	User or group name
type	string	path	True	User or group type

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
permission	string	<p>Specifies the access rights that a user or group has on the defined CIFS Share.</p> <p>The following values are allowed:</p> <ul style="list-style-type: none"> • no_access - User does not have CIFS share access • read - User has only read access • change - User has change access • full_control - User has full_control access
type	string	<p>Specifies the type of the user or group to add to the access control list of a CIFS share. The following values are allowed:</p> <ul style="list-style-type: none"> • windows - Windows user or group • unix_user - UNIX user • unix_group - UNIX group
user_or_group	string	<p>Specifies the user or group name to add to the access control list of a CIFS share.</p>

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "permission": "string",
  "type": "string",
  "user_or_group": "ENGDOMAIN\\ad_user"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a share-level ACL on a CIFS share

```
PATCH /protocols/cifs/shares/{svm.uuid}/{share}/acls/{user_or_group}/{type}
```

Introduced In: 9.6

Updates a share-level ACL on a CIFS share.

Related ONTAP commands

- `vserver cifs share access-control modify`

Learn more

- [DOC /protocols/cifs/shares/{svm.uuid}/{share}/acls](#)

Parameters

Name	Type	In	Required	Description
share	string	path	True	Share name
user_or_group	string	path	True	User or group name
type	string	path	True	User or group type
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
permission	string	<p>Specifies the access rights that a user or group has on the defined CIFS Share.</p> <p>The following values are allowed:</p> <ul style="list-style-type: none">• no_access - User does not have CIFS share access• read - User has only read access• change - User has change access• full_control - User has full_control access

Example request

```
{
  "permission": "string"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655516	The share ACL does not exist for given user and share

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cifs_share_acl

The permissions that users and groups have on a CIFS share.

Name	Type	Description
permission	string	<p>Specifies the access rights that a user or group has on the defined CIFS Share.</p> <p>The following values are allowed:</p> <ul style="list-style-type: none">• no_access - User does not have CIFS share access• read - User has only read access• change - User has change access• full_control - User has full_control access

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Manage CIFS UNIX symlink mapping

Protocols CIFS unix-symlink-mapping endpoint overview

Overview

ONTAP allows both CIFS and NFS to access the same datastore. This datastore can contain symbolic links which are files, created by UNIX clients. It contains a reference to another file or directory. If an SMB client accesses a symbolic link, it is redirected to the target file or directory that the symbolic link refers to. The symbolic links can point to files within the volume that contain the share, or to files that are contained in other volumes on the Storage Virtual Machine (SVM), or even to volumes contained on other SVMs.

There are two types of symbolic links:

Relative A relative symbolic link contains a reference to the file or directory relative to its parent directory. Therefore, the path of the file it is referring to should not begin with a backslash (/). If you enable symbolic links on a share, relative symbolic links work without UNIX symlink mapping.

Absolute An absolute symbolic link contains a reference to a file or directory in the form of an absolute path. Therefore, the path of the file it is referring to should begin with a backslash (/). An absolute symbolic link can refer to a file or directory within or outside of the file system of the symbolic link. If the target is not in the same local file system, the symbolic link is called a "widelink". If the symbolic link is enabled on a share and absolute symbolic links do not work right away, the mapping between the UNIX path of the symbolic link to the destination CIFS path must be created. When creating absolute symbolic link mappings, locality could be either "local" or "widelink" and it must be specified. If UNIX symlink mapping is created for a file or directory which is outside of the local share but the locality is set to "local", ONTAP does not allow access to the target.

A UNIX symbolic link support could be added to SMB shares by specifying the *unix_symlink* property during the creation of SMB shares or at any time by modifying the existing SMB *unix_symlink* property. UNIX symbolic link support is enabled by default.

Examples

Creating a UNIX symlink mapping for CIFS shares

To create UNIX symlink mappings for SMB shares, use the following API. Note the *return_records=true* query parameter used to obtain the newly created entry in the response.

```
# The API:
POST /api/protocols/cifs/unix-symlink-mapping

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-
mapping?return_records=true" -H "accept: application/json" -H "Content-
Type: application/json" -d "{ \"svm\": { \"name\": \"vs1\", \"uuid\":
\"000c5cd2-ebdf-11e8-a96e-0050568ea3cb\" }, \"target\": {
\"home_directory\": false, \"locality\": \"local\", \"path\":
\"/dir1/dir2/\", \"server\": \"cifs123\", \"share\": \"sh1\" },
\"unix_path\": \"/mnt/eng_volume/\"}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
        "name": "vs1"
      },
      "unix_path": "/mnt/eng_volume/",
      "target": {
        "share": "sh1",
        "path": "/dir1/dir2/",
        "server": "cifs123",
        "locality": "local",
        "home_directory": false
      }
    }
  ]
}
```

Retrieving UNIX symlink mappings for all SVMs in the cluster

```
# The API:
GET /api/protocols/cifs/unix-symlink-mapping

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-
mapping?fields=*&return_records=true&return_timeout=15" -H "accept:
application/hal+json"
```



```
# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/000c5cd2-ebdf-11e8-a96e-0050568ea3cb"
          }
        }
      },
      "unix_path": "/mnt/eng_volume/",
      "target": {
        "share": "sh1",
        "path": "/dir1/dir2/",
        "server": "CIFS123",
        "locality": "local",
        "home_directory": false
      },
      "_links": {
        "self": {
          "href": "/api/protocols/cifs/unix-symlink-mapping/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F"
        }
      }
    },
    {
      "svm": {
        "uuid": "1d30d1b1-ebdf-11e8-a96e-0050568ea3cb",
        "name": "vs2",
        "_links": {
          "self": {
            "href": "/api/svm/svms/1d30d1b1-ebdf-11e8-a96e-0050568ea3cb"
          }
        }
      },
      "unix_path": "/mnt/eng_volume/",
      "target": {
        "share": "ENG_SHARE",
        "path": "/dir1/dir2/",
        "server": "ENG_CIFS",
        "locality": "widelink",
        "home_directory": false
      }
    }
  ]
}
```

```

    "_links": {
      "self": {
        "href": "/api/protocols/cifs/unix-symlink-mapping/1d30d1b1-ebdf-11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F"
      }
    }
  },
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/protocols/cifs/unix-symlink-mapping?fields=*&return_records=true&return_timeout=15"
    }
  }
}

```

Retrieving a specific UNIX symlink mapping for an SVM

The mapping being returned is identified by the UUID of its SVM and the unix-path.

```

# The API:
GET /api/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-mapping/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "000c5cd2-ebdf-11e8-a96e-0050568ea3cb",
    "name": "vs1"
  },
  "unix_path": "/mnt/eng_volume/",
  "target": {
    "share": "sh1",
    "path": "/dir1/dir2/",
    "server": "CIFS123",
    "locality": "local",
    "home_directory": false
  }
}

```

Updating a specific UNIX symlink mapping for an SVM

The mapping being modified is identified by the UUID of its SVM and the unix-path.

```
# The API:
PATCH /api/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-mapping/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F" -H "accept: application/json" -H "Content-Type: application/json" -d '{"target\": { \"home_directory\": true, \"locality\": \"widelink\", \"path\": \"/new_path/\", \"server\": \"HR_SERVER\", \"share\": \"sh2\" } }'
```

Removing a specific UNIX symlink mapping for an SVM

The mapping being removed is identified by the UUID of its SVM and the unix-path.

```
# The API:
DELETE /api/protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/cifs/unix-symlink-mapping/000c5cd2-ebdf-11e8-a96e-0050568ea3cb/%2Fmnt%2Feng_volume%2F" -H "accept: application/json"
```

Retrieve UNIX symbolic link mappings for CIFS clients

GET /protocols/cifs/unix-symlink-mapping

Introduced In: 9.6

Retrieves UNIX symbolic link mappings for CIFS clients.

Related ONTAP commands

- `vserver cifs symlink show`

Learn more

- [DOC /protocols/cifs/unix-symlink-mapping](#)

Parameters

Name	Type	In	Required	Description
unix_path	string	query	False	Filter by unix_path
target.share	string	query	False	Filter by target.share
target.home_directory	boolean	query	False	Filter by target.home_directory
target.server	string	query	False	Filter by target.server
target.path	string	query	False	Filter by target.path
target.locality	string	query	False	Filter by target.locality
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cifs_symlink_mapping]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "locality": "string",
        "path": "/dir1/dir2/",
        "server": "ENGCIIFS",
        "share": "ENG_SHARE"
      },
      "unix_path": "/mnt/eng_volume/"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_target

Name	Type	Description
home_directory	boolean	Specify if the destination share is a home directory.
locality	string	<p>Specifies whether the CIFS symbolic link is a local link or wide link. The following values are supported:</p> <ul style="list-style-type: none">• local - Local symbolic link maps only to the same CIFS share.• widelink - Wide symbolic link maps to any CIFS share on the network.

Name	Type	Description
path	string	Specifies the CIFS path on the destination to which the symbolic link maps. The final path is generated by concatenating the CIFS server name, the share name, the cifs-path and the remaining path in the symbolic link left after the prefix match. This value is specified by using a UNIX-style path name. The trailing forward slash is required for the full path name to be properly interpreted.
server	string	Specifies the destination CIFS server where the UNIX symbolic link is pointing. This field is mandatory if the locality of the symbolic link is 'widelink'. You can specify the value in any of the following formats: <ul style="list-style-type: none"> • DNS name of the CIFS server. • IP address of the CIFS server. • NetBIOS name of the CIFS server.
share	string	Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing.

cifs_symlink_mapping

ONTAP allows for both CIFS and NFS access to the same datastore. This datastore can contain symbolic links created by UNIX clients which can point anywhere from the perspective of the UNIX client. To Access such UNIX symlink from CIFS share, we need to create a CIFS symbolic link path mapping from a UNIX symlink and target it as a CIFS path.

Name	Type	Description
_links	_links	
svm	svm	
target	cifs_target	

Name	Type	Description
unix_path	string	Specifies the UNIX path prefix to be matched for the mapping.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a UNIX symbolic link mapping for a CIFS client

POST /protocols/cifs/unix-symlink-mapping

Introduced In: 9.6

Creates a UNIX symbolic link mapping for a CIFS client.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the CIFS unix-symlink-mapping.
- `unix_path` - UNIX path to which the CIFS symlink mapping to be created.
- `target.share` - CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing.
- `target.path` - CIFS path on the destination to which the symbolic link maps.

Default property values

- `target.server` - *Local_NetBIOS_Server_Name*
- `locality` - *local*

- `home_directory` - *false*

Related ONTAP commands

- `vserver cifs symlink create`

Learn more

- [DOC /protocols/cifs/unix-symlink-mapping](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
svm	svm	
target	cifs_target	
unix_path	string	Specifies the UNIX path prefix to be matched for the mapping.

Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "locality": "string",
    "path": "/dir1/dir2/",
    "server": "ENG_CIFS",
    "share": "ENG_SHARE"
  },
  "unix_path": "/mnt/eng_volume/"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[cifs_symlink_mapping]	

Example response

```
{
  "records": [
    {
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "target": {
        "locality": "string",
        "path": "/dir1/dir2/",
        "server": "ENG_CIFS",
        "share": "ENG_SHARE"
      },
      "unix_path": "/mnt/eng_volume/"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655654	Must specify the target CIFS share while creating path mapping entries with localities "local" or "widelink"
655572	The target path contains illegal characters or is too long
655574	The target server contains illegal characters or is too long

Error Code	Description
655436	If the locality is "local", the target server must be blank or must match the CIFS NetBIOS name for given SVM
655439	The Specified target server is local CIFS server for given SVM but the locality is specified as "widelink"
655546	Failed to create symlink mapping because administrative share cannot be used as target share
655437	Failed to create the symlink mapping with locality "local" because the target share does not exist for specified SVM
655429	UNIX path must begin and end with a "/"
655430	Target path must begin and end with a "/"
655399	Failed to get the CIFS server for specified SVM

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_target

Name	Type	Description
home_directory	boolean	Specify if the destination share is a home directory.
locality	string	<p>Specifies whether the CIFS symbolic link is a local link or wide link.</p> <p>The following values are supported:</p> <ul style="list-style-type: none">• local - Local symbolic link maps only to the same CIFS share.• widelink - Wide symbolic link maps to any CIFS share on the network.
path	string	<p>Specifies the CIFS path on the destination to which the symbolic link maps. The final path is generated by concatenating the CIFS server name, the share name, the cifs-path and the remaining path in the symbolic link left after the prefix match. This value is specified by using a UNIX-style path name. The trailing forward slash is required for the full path name to be properly interpreted.</p>

Name	Type	Description
server	string	Specifies the destination CIFS server where the UNIX symbolic link is pointing. This field is mandatory if the locality of the symbolic link is 'widelink'. You can specify the value in any of the following formats: <ul style="list-style-type: none"> • DNS name of the CIFS server. • IP address of the CIFS server. • NetBIOS name of the CIFS server.
share	string	Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing.

cifs_symlink_mapping

ONTAP allows for both CIFS and NFS access to the same datastore. This datastore can contain symbolic links created by UNIX clients which can point anywhere from the perspective of the UNIX client. To Access such UNIX symlink from CIFS share, we need to create a CIFS symbolic link path mapping from a UNIX symlink and target it as a CIFS path.

Name	Type	Description
svm	svm	
target	cifs_target	
unix_path	string	Specifies the UNIX path prefix to be matched for the mapping.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete the UNIX symbolic link mapping for CIFS clients

```
DELETE /protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}
```

Introduced In: 9.6

Deletes the UNIX symbolic link mapping for CIFS clients.

Related ONTAP commands

- `vserver cifs symlink delete`

Learn more

- [DOC /protocols/cifs/unix-symlink-mapping](#)

Parameters

Name	Type	In	Required	Description
unix_path	string	path	True	UNIX symbolic link path
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Response

```
Status: 200, Ok
```

Error

```
Status: Default, Error
```


Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the UNIX symbolic link mapping for CIFS clients

GET /protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}

Introduced In: 9.6

Retrieves a UNIX symbolic link mapping for CIFS clients.

Related ONTAP commands

- `vserver cifs symlink show`

Learn more

- [DOC /protocols/cifs/unix-symlink-mapping](#)

Parameters

Name	Type	In	Required	Description
unix_path	string	path	True	UNIX symbolic link path
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
svm	svm	
target	cifs_target	
unix_path	string	Specifies the UNIX path prefix to be matched for the mapping.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "locality": "string",
    "path": "/dir1/dir2/",
    "server": "ENGCIFS",
    "share": "ENG_SHARE"
  },
  "unix_path": "/mnt/eng_volume/"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_target

Name	Type	Description
home_directory	boolean	Specify if the destination share is a home directory.
locality	string	<p>Specifies whether the CIFS symbolic link is a local link or wide link. The following values are supported:</p> <ul style="list-style-type: none">• local - Local symbolic link maps only to the same CIFS share.• widelink - Wide symbolic link maps to any CIFS share on the network.

Name	Type	Description
path	string	Specifies the CIFS path on the destination to which the symbolic link maps. The final path is generated by concatenating the CIFS server name, the share name, the cifs-path and the remaining path in the symbolic link left after the prefix match. This value is specified by using a UNIX-style path name. The trailing forward slash is required for the full path name to be properly interpreted.
server	string	Specifies the destination CIFS server where the UNIX symbolic link is pointing. This field is mandatory if the locality of the symbolic link is 'widelink'. You can specify the value in any of the following formats: <ul style="list-style-type: none"> • DNS name of the CIFS server. • IP address of the CIFS server. • NetBIOS name of the CIFS server.
share	string	Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the UNIX symbolic link mapping for CIFS clients

PATCH /protocols/cifs/unix-symlink-mapping/{svm.uuid}/{unix_path}

Introduced In: 9.6

Updates the UNIX symbolic link mapping for CIFS clients.

Related ONTAP commands

- `vserver cifs symlink modify`

Learn more

- [DOC /protocols/cifs/unix-symlink-mapping](#)

Parameters

Name	Type	In	Required	Description
unix_path	string	path	True	UNIX symbolic link path
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
svm	svm	
target	cifs_target	

Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "locality": "string",
    "path": "/dir1/dir2/",
    "server": "ENG_CIFS",
    "share": "ENG_SHARE"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
655573	Failed to modify the symlink mapping to target path because it contains illegal characters or is too long
655575	Failed to modify the symlink mapping to target server because it contains illegal characters or is too long
655547	Failed to modify symlink mapping because administrative share cannot be used as target share

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cifs_target

Name	Type	Description
home_directory	boolean	Specify if the destination share is a home directory.
locality	string	<p>Specifies whether the CIFS symbolic link is a local link or wide link.</p> <p>The following values are supported:</p> <ul style="list-style-type: none">• local - Local symbolic link maps only to the same CIFS share.• widelink - Wide symbolic link maps to any CIFS share on the network.
path	string	<p>Specifies the CIFS path on the destination to which the symbolic link maps. The final path is generated by concatenating the CIFS server name, the share name, the cifs-path and the remaining path in the symbolic link left after the prefix match. This value is specified by using a UNIX-style path name. The trailing forward slash is required for the full path name to be properly interpreted.</p>

Name	Type	Description
server	string	Specifies the destination CIFS server where the UNIX symbolic link is pointing. This field is mandatory if the locality of the symbolic link is 'widelink'. You can specify the value in any of the following formats: <ul style="list-style-type: none"> • DNS name of the CIFS server. • IP address of the CIFS server. • NetBIOS name of the CIFS server.
share	string	Specifies the CIFS share name on the destination CIFS server to which the UNIX symbolic link is pointing.

cifs_symlink_mapping

ONTAP allows for both CIFS and NFS access to the same datastore. This datastore can contain symbolic links created by UNIX clients which can point anywhere from the perspective of the UNIX client. To Access such UNIX symlink from CIFS share, we need to create a CIFS symbolic link path mapping from a UNIX symlink and target it as a CIFS path.

Name	Type	Description
svm	svm	
target	cifs_target	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve trace results for denied or allowed events

GET /protocols/file-access-tracing/events

Introduced In: 9.8

Retrieves the trace results for access allowed or denied events.

Related ONTAP commands

- `vserver security trace trace-result show`

Parameters

Name	Type	In	Required	Description
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
reason.message	string	query	False	Filter by reason.message
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
index	integer	query	False	Filter by index
session_id	integer	query	False	Filter by session_id
filter.trace_allowed_ops	boolean	query	False	Filter by filter.trace_allowed_ops
filter.unix_user	string	query	False	Filter by filter.unix_user

Name	Type	In	Required	Description
filter.index	integer	query	False	Filter by filter.index
filter.client_ip	string	query	False	Filter by filter.client_ip
filter.svm.uuid	string	query	False	Filter by filter.svm.uuid
filter.svm.name	string	query	False	Filter by filter.svm.name
filter.enabled	boolean	query	False	Filter by filter.enabled
filter.windows_user	string	query	False	Filter by filter.windows_user
filter.path	string	query	False	Filter by filter.path
filter.protocol	string	query	False	Filter by filter.protocol
share.name	string	query	False	Filter by share.name
create_time	string	query	False	Filter by create_time
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	collection_links	
num_records	integer	Number of records
records	array[file_access_event]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "create_time": "2018-06-04T19:00:00Z",
      "filter": {
        "client_ip": "10.140.68.143",
        "index": "1",
        "path": "/dir1/dir2",
        "protocol": "string",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "unix_user": "root",
        "windows_user": "cifs1/administrator"
      },
      "index": "1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
  ],
}
```

```

    "reason": {
      "message": "Access is allowed because the operation is trusted
and no security is configured."
    },
    "session_id": "2628976282477527056",
    "share": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "sh1"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

collection_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm_reference

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

file_access_filter

ONTAP allows creation of filters for file access tracing for both CIFS and NFS. These filters have protocols, path, username and client IP based on which file access operations are logged.

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.

Name	Type	Description
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.
svm	svm_reference	SVM, applies only to SVM-scoped objects.
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

reason

Displays the allowed or denied reason for the file access tracing events that are generated.

Name	Type	Description
message	string	The error message.

share

Name	Type	Description
_links	_links	
name	string	Share name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

file_access_event

ONTAP generates the list of file access tracing records stored on the cluster. These records are generated in response to security trace filters applied. The list of trace events recorded depends on the parameters configured for the filter.

Name	Type	Description
_links	_links	
create_time	string	Specifies the time at which the trace event entry was generated.
filter	file_access_filter	ONTAP allows creation of filters for file access tracing for both CIFS and NFS. These filters have protocols, path, username and client IP based on which file access operations are logged. <ul style="list-style-type: none">• Introduced in: 9.8
index	integer	Specifies the sequence number of the security trace event.
node	node	

Name	Type	Description
reason	reason	Displays the allowed or denied reason for the file access tracing events that are generated.
session_id	integer	Specifies the CIFS session ID for the file access trace event, this is generated only for CIFS file accesses.
share	share	
svm	svm_reference	SVM, applies only to SVM-scoped objects.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete trace results

DELETE /protocols/file-access-tracing/events/{node.uuid}/{svm.uuid}/{index}

Introduced In: 9.8

Deletes trace results.

Related ONTAP commands

- vserver security trace result delete

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID.
index	integer	path	True	Sequence number of traced event.
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve trace results

GET /protocols/file-access-tracing/events/{node.uuid}/{svm.uuid}/{index}

Introduced In: 9.8

Retrieves trace results for the specified sequence number.

Related ONTAP commands

- `vserver security trace trace-result show`

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID.
index	integer	path	True	Sequence number of traced event.

Name	Type	In	Required	Description
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
create_time	string	Specifies the time at which the trace event entry was generated.
filter	file_access_filter	<p>ONTAP allows creation of filters for file access tracing for both CIFS and NFS. These filters have protocols, path, username and client IP based on which file access operations are logged.</p> <ul style="list-style-type: none"> Introduced in: 9.8
index	integer	Specifies the sequence number of the security trace event.
node	node	
reason	reason	Displays the allowed or denied reason for the file access tracing events that are generated.
session_id	integer	Specifies the CIFS session ID for the file access trace event, this is generated only for CIFS file accesses.
share	share	
svm	svm_reference	SVM, applies only to SVM-scoped objects.
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "create_time": "2018-06-04T19:00:00Z",
  "filter": {
    "client_ip": "10.140.68.143",
    "index": "1",
    "path": "/dir1/dir2",
    "protocol": "string",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "unix_user": "root",
    "windows_user": "cifs1/administrator"
  },
  "index": "1",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "reason": {
    "message": "Access is allowed because the operation is trusted and no security is configured."
  },
  "session_id": "2628976282477527056",
  "share": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```



```

    },
    "name": "sh1"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm_reference

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

file_access_filter

ONTAP allows creation of filters for file access tracing for both CIFS and NFS. These filters have protocols, path, username and client IP based on which file access operations are logged.

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.

Name	Type	Description
protocol	string	Specifies the protocol for which permission trace is required.
svm	svm_reference	SVM, applies only to SVM-scoped objects.
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

reason

Displays the allowed or denied reason for the file access tracing events that are generated.

Name	Type	Description
message	string	The error message.

share

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	Share name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve security trace filter entry information

GET /protocols/file-access-tracing/filters

Introduced In: 9.8

Retrieves information about security trace filter entries.

Related ONTAP commands

- `vserver security trace filter show`

Parameters

Name	Type	In	Required	Description
trace_allowed_ops	boolean	query	False	Filter by trace_allowed_ops
unix_user	string	query	False	Filter by unix_user
index	integer	query	False	Filter by index
client_ip	string	query	False	Filter by client_ip
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
enabled	boolean	query	False	Filter by enabled
windows_user	string	query	False	Filter by windows_user
path	string	query	False	Filter by path
protocol	string	query	False	Filter by protocol
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	collection_links	
num_records	integer	Number of records
records	array[file_access_filter]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "client_ip": "10.140.68.143",
      "index": "1",
      "path": "/dir1/dir2",
      "protocol": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "unix_user": "root",
      "windows_user": "cifs1/administrator"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

collection_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm_reference

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

file_access_filter

ONTAP allows creation of filters for file access tracing for both CIFS and NFS. These filters have protocols, path, username and client IP based on which file access operations are logged.

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.

Name	Type	Description
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.
svm	svm_reference	SVM, applies only to SVM-scoped objects.
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Create security trace filter entries

POST /protocols/file-access-tracing/filters

Introduced In: 9.8

Creates security trace filter entries.

Related ONTAP commands

- `vserver security trace filter create`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.

Name	Type	Description
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.
svm	svm_reference	SVM, applies only to SVM-scoped objects.
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

Example request

```
{
  "client_ip": "10.140.68.143",
  "index": "1",
  "path": "/dir1/dir2",
  "protocol": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_user": "root",
  "windows_user": "cifs1/administrator"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[file_access_filter]	

Example response

```
{
  "records": [
    {
      "client_ip": "10.140.68.143",
      "index": "1",
      "path": "/dir1/dir2",
      "protocol": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "unix_user": "root",
      "windows_user": "cifs1/administrator"
    }
  ]
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm_reference

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

file_access_filter

ONTAP allows creation of filters for file access tracing for both CIFS and NFS. These filters have protocols, path, username and client IP based on which file access operations are logged.

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.
svm	svm_reference	SVM, applies only to SVM-scoped objects.

Name	Type	Description
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

collection_links

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update security trace filter entries

PATCH /protocols/file-access-tracing/filters/{svm.uuid}/{index}

Introduced In: 9.8

Updates security trace filter entries.

Related ONTAP commands

- `vserver security trace filter modify`

Parameters

Name	Type	In	Required	Description
index	integer	path	True	Filter index.
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.

Name	Type	Description
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

Example request

```
{
  "client_ip": "10.140.68.143",
  "index": "1",
  "path": "/dir1/dir2",
  "protocol": "string",
  "unix_user": "root",
  "windows_user": "cifs1/administrator"
}
```

Response

Status: 200, Ok

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm_reference

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

file_access_filter

ONTAP allows creation of filters for file access tracing for both CIFS and NFS. These filters have protocols, path, username and client IP based on which file access operations are logged.

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.

Name	Type	Description
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve security trace filter entry information

GET /protocols/file-access-tracing/filters/{svm.uuid}/{index}

Introduced In: 9.8

Retrieves information about security trace filter entries.

Related ONTAP commands

- `vserver security trace filter show`

Parameters

Name	Type	In	Required	Description
index	integer	path	True	
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.
svm	svm_reference	SVM, applies only to SVM-scoped objects.

Name	Type	Description
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

Example response

```
{
  "client_ip": "10.140.68.143",
  "index": "1",
  "path": "/dir1/dir2",
  "protocol": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_user": "root",
  "windows_user": "cifs1/administrator"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm_reference

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update security trace filter entries

PATCH /protocols/file-access-tracing/filters/{svm.uuid}/{index}

Introduced In: 9.8

Updates security trace filter entries.

Related ONTAP commands

- `vserver security trace filter modify`

Parameters

Name	Type	In	Required	Description
index	integer	path	True	Filter index.
svm.uuid	string	path	True	UUID of the SVM to which this object belongs.

Request Body

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.

Name	Type	Description
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

Example request

```
{
  "client_ip": "10.140.68.143",
  "index": "1",
  "path": "/dir1/dir2",
  "protocol": "string",
  "unix_user": "root",
  "windows_user": "cifs1/administrator"
}
```

Response

Status: 200, Ok

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm_reference

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

file_access_filter

ONTAP allows creation of filters for file access tracing for both CIFS and NFS. These filters have protocols, path, username and client IP based on which file access operations are logged.

Name	Type	Description
client_ip	string	Specifies the IP address from which the client accesses the file or directory.
enabled	boolean	Specifies whether to enable or disable the filter. Filters are enabled by default and are deleted after 60 mins.
index	integer	Position of the file access tracing filter.
path	string	Specifies the path for which permission tracing can be applied. The value can be complete path from root of CIFS share or root of volume for NFS.
protocol	string	Specifies the protocol for which permission trace is required.

Name	Type	Description
trace_allowed_ops	boolean	Specifies if the filter can trace file access denied and allowed events. The value of trace-allow is false by default, and it traces access denied events. The value is set to true for tracing access allowed events.
unix_user	string	Specifies the UNIX username whose access requests you want to trace. The filter would match only if the request is received with this user.
windows_user	string	Specifies the Windows username whose access requests you want to trace. The filter would match only if the request is received with this user.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View file security permissions

Protocols file-security effective-permissions svm.uuid path endpoint overview

Overview

This API displays the effective permission granted to a Windows or UNIX user on the specified file or folder path.

Examples

Retrieving the effective permission for the specified Windows user on the specified path of an SVM.

= The API:

```
curl -X GET "https://10.63.26.252/api/protocols/file-security/effective-permissions/cf5f271a-1beb-11ea-8fad-005056bb645e/administrator/windows/%2F?share.name=sh1&return_records=true" -H "accept: application/json" -H "authorization: Basic YWRtaW46bmV0YXBwMSE="
```

= The response:

```
{
  "svm": {
    "uuid": "cf5f271a-1beb-11ea-8fad-005056bb645e",
    "name": "vs1"
  },
  "user": "administrator",
  "type": "windows",
  "path": "/",
  "share": {
    "path": "/"
  },
  "file_permission": [
    "read",
    "write",
    "append",
    "read_ea",
    "write_ea",
    "execute",
    "delete_child",
    "read_attributes",
    "write_attributes",
    "delete",
    "read_control",
    "write_dac",
    "write_owner",
    "synchronize",
    "system_security"
  ],
}
```

```
"share_permission": [
  "read",
  "read_ea",
  "execute",
  "read_attributes",
  "read_control",
  "synchronize"
]
}
```

```
[[IDbelf5d6754b5b17a9bc358ad075e9158]]
= Retrieve effective security permissions on a file
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/file-security/effective-permissions/{svm.uuid}/{path}`#
```

Introduced In: 9.8

Retrieves effective security permissions on a file.

== Related ONTAP commands

* `vserver security file-directory show-effective-permissions`

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|path
|string
|path
|True
a|File Path
```

```
|user
|string
|query
|True
a|User_Name

|share.name
|string
|query
|False
a|Share Name

|type
|string
|query
|False
a|User Type

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```



```

|file_permissions
|array[string]
a|Specifies the effective permission granted to a user on the specified
file or folder path.

|path
|string
a|Specifies the path of the file or the folder for which you want to
display effective permissions.
The path is relative to the SVM root volume. If "-share-name" is specified
then path will be relative to the share path.

|share
|link:#share[share]
a|

|share_permissions
|array[string]
a|Specifies the effective permission granted to a user on the specified
file or folder path.

|svm
|link:#svm_reference[svm_reference]
a|SVM, applies only to SVM-scoped objects.

|type
|string
a|Specifies the user type. The following values are allowed:

* windows    - Windows user
* unix       - UNIX user

|user
|string
a|Specifies the user for which effective permission needs to be displayed
for the specified path.

|===

```

.Example response

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "file_permissions": [
    "string"
  ],
  "path": "/dir1/dir2",
  "share": {
    "name": "string",
    "path": "string"
  },
  "share_permissions": [
    "string"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "user": "cifs1/administrator"
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#share]
[.api-collapsible-fifth-title]
share

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Displays the file or directory effective permission for the mentioned
user, only for files and directories contained where the
specified path is relative to the root of the specified share. If this
parameter is not specified, the SVM root volume is
taken as the default. If this parameter is specified, the effective share
permission of the user is also displayed.
Wildcard query characters are not supported.
```

```
|path
|string
a|Displays the CIFS share path.
```

```
|===
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#svm_reference]
[.api-collapsible-fifth-title]
svm_reference
```

SVM, applies only to SVM-scoped objects.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage FPolicy configuration

:leveloffset: +1

```

```
[[ID4b6af50af9764d5fd5afe3a270b08808]]
```

= Protocols fpolicy endpoint overview

== Overview

FPolicy is an infrastructure component of ONTAP that enables partner applications to connect to ONTAP in order to monitor and set file access permissions. Every time a client accesses a file from a storage system, based on the configuration of FPolicy, the partner application is notified about file access. This enables partners to set restrictions on files that are created or accessed on the storage system. FPolicy also allows you to create file policies that specify file operation permissions according to file type. For example, you can restrict certain file types, such as .jpeg and .mp3 files, from being stored on the storage system. FPolicy can monitor file access from CIFS and NFS clients.

As part of FPolicy configuration, you can specify an FPolicy engine which defines the external FPolicy server, FPolicy events, which defines the protocol and file operations to monitor and the FPolicy policy that acts as a container for the FPolicy engine and FPolicy events. It provides a way for policy management functions, such as policy enabling and disabling.

== Examples

=== Creating an FPolicy configuration

To create an FPolicy for an SVM use the following API. Note that the `_return_records=true_` query parameter is used to obtain the newly created entry in the response.

The API:

POST /protocols/fpolicy/

#The call:

```
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"engines\": [ { \"name\": \"engine1\", \"port\": 9876,
  \"primary_servers\": [ \"10.132.145.22\", \"10.140.101.109\" ],
  \"secondary_servers\": [ \"10.132.145.20\", \"10.132.145.21\" ], \"type\":
  \"synchronous\" } ], \"events\": [ { \"file_operations\": { \"read\":
  true, \"write\": true }, \"filters\": { \"monitor_ads\": true }, \"name\":
  \"event_cifs\", \"protocol\": \"cifs\", \"volume_monitoring\": true } ],
```

```
\\"policies\\": [ { \\"engine\\": { \\"name\\": \\"engine1\\" }, \\"events\\": [ {  
  \\"name\\": \\"event_cifs\\" } ], \\"mandatory\\": true, \\"name\\": \\"pol0\\",  
  \\"priority\\": 1, \\"scope\\": { \\"include_volumes\\": [ \\"vol1\\" ] } } ],  
  \\"svm\\": { \\"name\\": \\"vs1\\", \\"uuid\\": \\"b34f5e3d-01d0-11e9-8f63-  
0050568ea311\\" } } ]"
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "svm": {  
        "uuid": "b34f5e3d-01d0-11e9-8f63-0050568ea311",  
        "name": "vs1"  
      },  
      "engines": [  
        {  
          "name": "engine1",  
          "primary_servers": [  
            "10.132.145.22",  
            "10.140.101.109"  
          ],  
          "secondary_servers": [  
            "10.132.145.20",  
            "10.132.145.21"  
          ],  
          "type": "synchronous",  
          "port": 9876  
        }  
      ],  
      "events": [  
        {  
          "name": "event_cifs",  
          "protocol": "cifs",  
          "volume_monitoring": true,  
          "file_operations": {  
            "read": true,  
            "write": true  
          },  
          "filters": {  
            "monitor_ads": true  
          }  
        }  
      ],  
      "policies": [  
        {
```



```

        "name": "pol0",
        "priority": 1,
        "events": [
            {
                "name": "event_cifs"
            }
        ],
        "engine": {
            "name": "engine1"
        },
        "scope": {
            "include_volumes": [
                "vol1"
            ]
        },
        "mandatory": true
    }
]
}
}

```

'''

=== Retrieving the FPolicy configuration for all the SVMs in the cluster

'''

The API:

GET /protocols/fpolicy

The call:

```

curl -X GET "https://<mgmt-
ip>/api/protocols/fpolicy?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"

```

The response:

```

{
  "records": [
    {
      "svm": {
        "uuid": "b34f5e3d-01d0-11e9-8f63-0050568ea311",
        "name": "vs1"
      }
    }
  ]
}

```

```

},
"engines": [
  {
    "name": "engine1",
    "primary_servers": [
      "10.132.145.22",
      "10.140.101.109"
    ],
    "secondary_servers": [
      "10.132.145.20",
      "10.132.145.21"
    ],
    "type": "synchronous",
    "port": 9876
  }
],
"events": [
  {
    "name": "event_cifs",
    "protocol": "cifs",
    "volume_monitoring": true,
    "file_operations": {
      "close": false,
      "create": false,
      "create_dir": false,
      "delete": false,
      "delete_dir": false,
      "getattr": false,
      "link": false,
      "lookup": false,
      "open": false,
      "read": true,
      "write": true,
      "rename": false,
      "rename_dir": false,
      "setattr": false,
      "symlink": false
    },
    "filters": {
      "monitor_ads": true,
      "close_with_modification": false,
      "close_without_modification": false,
      "close_with_read": false,
      "first_read": false,
      "first_write": false,
      "offline_bit": false,

```

```

        "open_with_delete_intent": false,
        "open_with_write_intent": false,
        "write_with_size_change": false,
        "setattr_with_owner_change": false,
        "setattr_with_group_change": false,
        "setattr_with_sacl_change": false,
        "setattr_with_dacl_change": false,
        "setattr_with_modify_time_change": false,
        "setattr_with_access_time_change": false,
        "setattr_with_creation_time_change": false,
        "setattr_with_mode_change": false,
        "setattr_with_size_change": false,
        "setattr_with_allocation_size_change": false,
        "exclude_directory": false
    }
}
],
"policies": [
    {
        "name": "pol0",
        "enabled": true,
        "priority": 1,
        "events": [
            {
                "name": "event_cifs"
            }
        ],
        "engine": {
            "name": "engine1"
        },
        "scope": {
            "include_volumes": [
                "vol1"
            ]
        },
        "mandatory": true
    }
]
}
],
"num_records": 1
}
----

'''

```

```
=== Retrieving an FPolicy configuration for a particular SVM
```

```
'''
```

```
----
```

```
# The API:
```

```
GET /protocols/fpolicy/{svm.uuid}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/b34f5e3d-01d0-11e9-8f63-0050568ea311?fields=*&return_records=true&return_timeout=15" -H "accept: application/json"
```

```
# The response:
```

```
{
  "svm": {
    "uuid": "b34f5e3d-01d0-11e9-8f63-0050568ea311",
    "name": "vs1"
  },
  "engines": [
    {
      "name": "engine1",
      "primary_servers": [
        "10.132.145.22",
        "10.140.101.109"
      ],
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "type": "synchronous",
      "port": 9876
    }
  ],
  "events": [
    {
      "name": "event_cifs",
      "protocol": "cifs",
      "volume_monitoring": true,
      "file_operations": {
        "close": false,
        "create": false,
        "create_dir": false,
        "delete": false,
        "delete_dir": false,

```

```

    "getattr": false,
    "link": false,
    "lookup": false,
    "open": false,
    "read": true,
    "write": true,
    "rename": false,
    "rename_dir": false,
    "setattr": false,
    "symlink": false
  },
  "filters": {
    "monitor_ads": true,
    "close_with_modification": false,
    "close_without_modification": false,
    "close_with_read": false,
    "first_read": false,
    "first_write": false,
    "offline_bit": false,
    "open_with_delete_intent": false,
    "open_with_write_intent": false,
    "write_with_size_change": false,
    "setattr_with_owner_change": false,
    "setattr_with_group_change": false,
    "setattr_with_sacl_change": false,
    "setattr_with_dacl_change": false,
    "setattr_with_modify_time_change": false,
    "setattr_with_access_time_change": false,
    "setattr_with_creation_time_change": false,
    "setattr_with_mode_change": false,
    "setattr_with_size_change": false,
    "setattr_with_allocation_size_change": false,
    "exclude_directory": false
  }
},
"policies": [
  {
    "name": "pol0",
    "enabled": true,
    "priority": 1,
    "events": [
      {
        "name": "event_cifs"
      }
    ]
  }
],

```

```

    "engine": {
        "name": "engine1"
    },
    "scope": {
        "include_volumes": [
            "vol1"
        ]
    },
    "mandatory": true
}
]
}
-----

'''

=== Deleting an FPolicy configuration for a particular SVM

'''

-----

# The API:
DELETE /protocols/fpolicy/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/fpolicy/b34f5e3d-01d0-
11e9-8f63-0050568ea311" -H "accept: application/json"
-----

'''

[[IDbf19dbfd5bd03821a1797acc5e667336]]
= Retrieve an FPolicy configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/fpolicy`#

*Introduced In:* 9.6

Retrieves an FPolicy configuration.

== Related ONTAP commands

```

- * ``fpolicy show``
- * ``fpolicy policy show``
- * ``fpolicy policy scope show``
- * ``fpolicy policy event show``
- * ``fpolicy policy external-engine show``

== Learn more

* `xref:{relative_path}protocols_fpolicy_endpoint_overview.html` [DOC
/protocols/fpolicy]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|engines.secondary_servers

|string

|query

|False

a|Filter by engines.secondary_servers

|engines.port

|integer

|query

|False

a|Filter by engines.port

|engines.type

|string

|query

|False

a|Filter by engines.type

|engines.name

|string

```
|query
|False
a|Filter by engines.name

|engines.primary_servers
|string
|query
|False
a|Filter by engines.primary_servers

|policies.priority
|integer
|query
|False
a|Filter by policies.priority

|policies.engine.name
|string
|query
|False
a|Filter by policies.engine.name

|policies.events.name
|string
|query
|False
a|Filter by policies.events.name

|policies.scope.include_shares
|string
|query
|False
a|Filter by policies.scope.include_shares

|policies.scope.include_extension
|string
|query
|False
a|Filter by policies.scope.include_extension
```



```
|policies.scope.include_volumes  
|string  
|query  
|False  
a|Filter by policies.scope.include_volumes
```

```
|policies.scope.include_export_policies  
|string  
|query  
|False  
a|Filter by policies.scope.include_export_policies
```

```
|policies.scope.exclude_export_policies  
|string  
|query  
|False  
a|Filter by policies.scope.exclude_export_policies
```

```
|policies.scope.exclude_shares  
|string  
|query  
|False  
a|Filter by policies.scope.exclude_shares
```

```
|policies.scope.exclude_volumes  
|string  
|query  
|False  
a|Filter by policies.scope.exclude_volumes
```

```
|policies.scope.exclude_extension  
|string  
|query  
|False  
a|Filter by policies.scope.exclude_extension
```

```
|policies.name  
|string  
|query  
|False  
a|Filter by policies.name
```

```
|policies.enabled  
|boolean  
|query  
|False  
a|Filter by policies.enabled
```

```
|policies.mandatory  
|boolean  
|query  
|False  
a|Filter by policies.mandatory
```

```
|events.protocol  
|string  
|query  
|False  
a|Filter by events.protocol
```

```
|events.name  
|string  
|query  
|False  
a|Filter by events.name
```

```
|events.filters.close_with_modification  
|boolean  
|query  
|False  
a|Filter by events.filters.close_with_modification
```

```
|events.filters.close_without_modification  
|boolean  
|query  
|False  
a|Filter by events.filters.close_without_modification
```

```
|events.filters.offline_bit  
|boolean  
|query
```

```

|False
a|Filter by events.filters.offline_bit

|events.filters.write_with_size_change
|boolean
|query
|False
a|Filter by events.filters.write_with_size_change

|events.filters.setattr_with_creation_time_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_creation_time_change

|events.filters.setattr_with_size_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_size_change

|events.filters.open_with_delete_intent
|boolean
|query
|False
a|Filter by events.filters.open_with_delete_intent

|events.filters.setattr_with_owner_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_owner_change

|events.filters.setattr_with_access_time_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_access_time_change

|events.filters.setattr_with_group_change

```

```
|boolean
|query
|False
a|Filter by events.filters.setattr_with_group_change
```

```
|events.filters.setattr_with_dacl_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_dacl_change
```

```
|events.filters.setattr_with_allocation_size_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_allocation_size_change
```

```
|events.filters.exclude_directory
|boolean
|query
|False
a|Filter by events.filters.exclude_directory
```

```
|events.filters.open_with_write_intent
|boolean
|query
|False
a|Filter by events.filters.open_with_write_intent
```

```
|events.filters.monitor_ads
|boolean
|query
|False
a|Filter by events.filters.monitor_ads
```

```
|events.filters.setattr_with_mode_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_mode_change
```

```
|events.filters.setattr_with_sacl_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_sacl_change
```

```
|events.filters.setattr_with_modify_time_change
|boolean
|query
|False
a|Filter by events.filters.setattr_with_modify_time_change
```

```
|events.filters.first_read
|boolean
|query
|False
a|Filter by events.filters.first_read
```

```
|events.filters.close_with_read
|boolean
|query
|False
a|Filter by events.filters.close_with_read
```

```
|events.filters.first_write
|boolean
|query
|False
a|Filter by events.filters.first_write
```

```
|events.volume_monitoring
|boolean
|query
|False
a|Filter by events.volume_monitoring
```

```
|events.file_operations.close
|boolean
|query
|False
```

```
a|Filter by events.file_operations.close

|events.file_operations.lookup
|boolean
|query
|False
a|Filter by events.file_operations.lookup

|events.file_operations.delete_dir
|boolean
|query
|False
a|Filter by events.file_operations.delete_dir

|events.file_operations.create_dir
|boolean
|query
|False
a|Filter by events.file_operations.create_dir

|events.file_operations.create
|boolean
|query
|False
a|Filter by events.file_operations.create

|events.file_operations.setattr
|boolean
|query
|False
a|Filter by events.file_operations.setattr

|events.file_operations.getattr
|boolean
|query
|False
a|Filter by events.file_operations.getattr

|events.file_operations.delete
|boolean
```

```
|query
|False
a|Filter by events.file_operations.delete

|events.file_operations.symlink
|boolean
|query
|False
a|Filter by events.file_operations.symlink

|events.file_operations.rename
|boolean
|query
|False
a|Filter by events.file_operations.rename

|events.file_operations.open
|boolean
|query
|False
a|Filter by events.file_operations.open

|events.file_operations.read
|boolean
|query
|False
a|Filter by events.file_operations.read

|events.file_operations.rename_dir
|boolean
|query
|False
a|Filter by events.file_operations.rename_dir

|events.file_operations.link
|boolean
|query
|False
a|Filter by events.file_operations.link
```

```
|events.file_operations.write
|boolean
|query
|False
a|Filter by events.file_operations.write
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
```



```

|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#fpolicy[fpolicy]]
a|

|===

```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "engines": [
        {
          "name": "fp_ex_eng",
          "port": "9876",
          "primary_servers": [
            "10.132.145.20",
            "10.140.101.109"
          ],
          "secondary_servers": [
            "10.132.145.20",
            "10.132.145.21"
          ],
          "type": "string"
        }
      ],
      "events": [
        {
          "name": "event_nfs_close",
          "protocol": "string"
        }
      ],
      "policies": [
        {
          "engine": {
            "_links": {
```

```

        "self": {
            "href": "/api/resourcelink"
        },
        "name": "string"
    },
    "events": [
        "event_nfs_close",
        "event_open"
    ],
    "name": "fp_policy_1",
    "scope": {
        "exclude_export_policies": [
            "string"
        ],
        "exclude_extension": [
            "string"
        ],
        "exclude_shares": [
            "string"
        ],
        "exclude_volumes": [
            "vol1",
            "vol_svm1",
            "*"
        ],
        "include_export_policies": [
            "string"
        ],
        "include_extension": [
            "string"
        ],
        "include_shares": [
            "sh1",
            "share_cifs"
        ],
        "include_volumes": [
            "vol1",
            "vol_svm1"
        ]
    }
}
],
"svm": {
    "_links": {
        "self": {

```

```

        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

```

====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#fpolicy_engine]
[.api-collapsible-fifth-title]
fpolicy_engine
```

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|Specifies the name to assign to the external server configuration.
```

```
|port
|integer
a|Port number of the FPolicy server application.
```

```
|primary_servers
|array[string]
a|
```

```
|secondary_servers
|array[string]
```

a|

|type

|string

a|The notification mode determines what ONTAP does after sending notifications to FPolicy servers.

The possible values are:

***** synchronous - After sending a notification, wait for a response from the FPolicy server.

***** asynchronous - After sending a notification, file request processing continues.

* Default value: 1

* enum: ["synchronous", "asynchronous"]

* Introduced in: 9.6

|===

[#file_operations]

[.api-collapsible-fifth-title]

file_operations

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.

The event will check the operations specified from all client requests using the protocol.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|close

|boolean

a|File close operations

|create

|boolean

a|File create operations

```
|create_dir  
|boolean  
a|Directory create operations
```

```
|delete  
|boolean  
a|File delete operations
```

```
|delete_dir  
|boolean  
a|Directory delete operations
```

```
|getattr  
|boolean  
a|Get attribute operations
```

```
|link  
|boolean  
a|Link operations
```

```
|lookup  
|boolean  
a|Lookup operations
```

```
|open  
|boolean  
a|File open operations
```

```
|read  
|boolean  
a|File read operations
```

```
|rename  
|boolean  
a|File rename operations
```

```
|rename_dir
```



```
|boolean  
a|Directory rename operations
```

```
|setattr  
|boolean  
a|Set attribute operations
```

```
|symlink  
|boolean  
a|Symbolic link operations
```

```
|write  
|boolean  
a|File write operations
```

```
|===
```

```
[#filters]  
[.api-collapsible-fifth-title]  
filters
```

Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|close_with_modification  
|boolean  
a|Filter the client request for close with modification.
```

```
|close_with_read  
|boolean  
a|Filter the client request for close with read.
```

|close_without_modification
|boolean
a|Filter the client request for close without modification.

|exclude_directory
|boolean
a|Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.

|first_read
|boolean
a|Filter the client requests for the first-read.

|first_write
|boolean
a|Filter the client requests for the first-write.

|monitor_ads
|boolean
a|Filter the client request for alternate data stream.

|offline_bit
|boolean
a|Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.

|open_with_delete_intent
|boolean
a|Filter the client request for open with delete intent.

|open_with_write_intent
|boolean
a|Filter the client request for open with write intent.

|setattr_with_access_time_change
|boolean
a|Filter the client setattr requests for changing the access time of a file or directory.

|setattr_with_allocation_size_change
|boolean
a|Filter the client setattr requests for changing the allocation size of a file.

|setattr_with_creation_time_change
|boolean
a|Filter the client setattr requests for changing the creation time of a file or directory.

|setattr_with_dacl_change
|boolean
a|Filter the client setattr requests for changing dacl on a file or directory.

|setattr_with_group_change
|boolean
a|Filter the client setattr requests for changing group of a file or directory.

|setattr_with_mode_change
|boolean
a|Filter the client setattr requests for changing the mode bits on a file or directory.

|setattr_with_modify_time_change
|boolean
a|Filter the client setattr requests for changing the modification time of a file or directory.

|setattr_with_owner_change
|boolean
a|Filter the client setattr requests for changing owner of a file or directory.

|setattr_with_sacl_change
|boolean
a|Filter the client setattr requests for changing sacl on a file or

directory.

|setattr_with_size_change

|boolean

a|Filter the client setattr requests for changing the size of a file.

|write_with_size_change

|boolean

a|Filter the client request for write with size change.

|===

[#fpolicy_event]

[.api-collapsible-fifth-title]

fpolicy_event

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|file_operations

|link:#file_operations[file_operations]

a|Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.

The event will check the operations specified from all client requests using the protocol.

|filters

|link:#filters[filters]

a|Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

```

|name
|string
a|Specifies the name of the FPolicy event.

|protocol
|string
a|Protocol for which event is created. If you specify protocol, then you
must also specify a valid value for the file operation parameters.
The value of this parameter must be one of the following:

***** cifs - for the CIFS protocol.

***** nfsv3 - for the NFSv3 protocol.

***** nfsv4 - for the NFSv4 protocol.

|volume_monitoring
|boolean
a|Specifies whether volume operation monitoring is required.

|===

[#fpolicy_engine_reference]
[.api-collapsible-fifth-title]
fpolicy_engine_reference

FPolicy external engine

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the FPolicy external engine.

```

```
|===
```

```
[#fpolicy_event_reference]  
[.api-collapsible-fifth-title]  
fpolicy_event_reference
```

FPolicy events

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|===
```

```
[#scope]
```

```
[.api-collapsible-fifth-title]  
scope
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|exclude_export_policies
```

```
|array[string]
```

```
a|
```

```
|exclude_extension
```

```
|array[string]
```

```
a|
```

```
|exclude_shares
```

```

|array[string]
a|

|exclude_volumes
|array[string]
a|

|include_export_policies
|array[string]
a|

|include_extension
|array[string]
a|

|include_shares
|array[string]
a|

|include_volumes
|array[string]
a|

|===

[#fpolicy_policy]
[.api-collapsible-fifth-title]
fpolicy_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies if the policy is enabled on the SVM or not. If no value is
mentioned for this field but priority is set, then this policy will be
enabled.

|engine
|link:#fpolicy_engine_reference[fpolicy_engine_reference]
a|FPolicy external engine

```

```

|events
|array[link:#fpolicy_event_reference[fpolicy_event_reference]]
a|

|mandatory
|boolean
a|Specifies what action to take on a file access event in a case when all
primary and secondary servers are down or no response is received from the
FPolicy servers within a given timeout period. When this parameter is set
to true, file access events will be denied under these circumstances.

|name
|string
a|Specifies the name of the policy.

|priority
|integer
a|Specifies the priority that is assigned to this policy.

|scope
|link:#scope[scope]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

```



```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#fpolicy]
[.api-collapsible-fifth-title]
fpolicy
```

FPolicy is an infrastructure component of ONTAP that enables partner applications connected to your storage systems to monitor and set file access permissions. Every time a client accesses a file from a storage system, based on the configuration of FPolicy, the partner application is notified about file access.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|engines
|array[link:#fpolicy_engine[fpolicy_engine]]
a|
```

```
|events
|array[link:#fpolicy_event[fpolicy_event]]
a|
```

```
|policies
|array[link:#fpolicy_policy[fpolicy_policy]]
a|
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
[#error_arguments]  
[.api-collapsible-fifth-title]  
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|arguments  
|array[link:#error_arguments[error_arguments]]  
a|Message arguments
```

```
|code  
|string  
a|Error code
```

```
|message
```

```

|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID68e2d8fcad2e7d177e1f0b978ca44eb1]]
= Create an FPolicy configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/fpolicy`#

*Introduced In:* 9.6

Creates an FPolicy configuration.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the FPolicy
configuration.

== Recommended optional properties

* `engines` - External server to which the notifications will be sent.
* `events` - File operations to monitor.
* `policies` - Policy configuration which acts as a container for FPolicy
event and FPolicy engine.
* `scope` - Scope of the policy. Can be limited to exports, volumes,
shares or file extensions.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `engines.type` - _synchronous_

```

```
* `policies.engine` - _native_
* `policies.mandatory` - _true_
* `events.volume_monitoring` - _false_
* `events.file_operations.+++` - _false_
* `events.filters.+++` - _false_
```

== Related ONTAP commands

```
* `fpolicy policy event create`
* `fpolicy policy external-engine create`
* `fpolicy policy create`
* `fpolicy policy scope create`
* `fpolicy enable`
```

== Learn more

```
* xref:{relative_path}protocols_fpolicy_endpoint_overview.html[DOC
/protocols/fpolicy]
```

== Parameters

```
[cols=5*,options=header]
|==
```

```
|Name
|Type
|In
|Required
|Description
```

```
|return_records
|boolean
|query
|False
```

```
a|The default is false. If set to true, the records are returned.
```

```
* Default value:
```

```
|==
```

== Request Body

```
[cols=3*,options=header]
|==
```

```

|Name
|Type
|Description

|engines
|array[link:#fpolicy_engine[fpolicy_engine]]
a|

|events
|array[link:#fpolicy_event[fpolicy_event]]
a|

|policies
|array[link:#fpolicy_policy[fpolicy_policy]]
a|

|svm
|link:#svm[svm]
a|

|===

```

.Example request

[%collapsible%closed]

====

```

[source,json,subs=+macros]
{
  "engines": [
    {
      "name": "fp_ex_eng",
      "port": "9876",
      "primary_servers": [
        "10.132.145.20",
        "10.140.101.109"
      ],
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "type": "string"
    }
  ],
  "events": [
    {
      "name": "event_nfs_close",

```

```

        "protocol": "string"
    }
],
"policies": [
{
    "enabled": null,
    "engine": {
        "name": "string"
    },
    "events": [
        "event_nfs_close",
        "event_open"
    ],
    "name": "fp_policy_1",
    "scope": {
        "exclude_export_policies": [
            "string"
        ],
        "exclude_extension": [
            "string"
        ],
        "exclude_shares": [
            "string"
        ],
        "exclude_volumes": [
            "vol1",
            "vol_svm1",
            "*"
        ],
        "include_export_policies": [
            "string"
        ],
        "include_extension": [
            "string"
        ],
        "include_shares": [
            "sh1",
            "share_cifs"
        ],
        "include_volumes": [
            "vol1",
            "vol_svm1"
        ]
    }
}
],

```

```

"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records

|records
|array[link:#fpolicy[fpolicy]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "engines": [
        {
          "name": "fp_ex_eng",
          "port": "9876",
          "primary_servers": [
            "10.132.145.20",
            "10.140.101.109"
          ],
          "secondary_servers": [
            "10.132.145.20",

```

```

        "10.132.145.21"
    ],
    "type": "string"
}
],
"events": [
    {
        "name": "event_nfs_close",
        "protocol": "string"
    }
],
"policies": [
    {
        "enabled": null,
        "engine": {
            "name": "string"
        },
        "events": [
            "event_nfs_close",
            "event_open"
        ],
        "name": "fp_policy_1",
        "scope": {
            "exclude_export_policies": [
                "string"
            ],
            "exclude_extension": [
                "string"
            ],
            "exclude_shares": [
                "string"
            ],
            "exclude_volumes": [
                "vol1",
                "vol_svm1",
                "*"
            ],
            "include_export_policies": [
                "string"
            ],
            "include_extension": [
                "string"
            ],
            "include_shares": [
                "sh1",
                "share_cifs"
            ]
        }
    }
]

```



```

        ],
        "include_volumes": [
            "vol1",
            "vol_svm1"
        ]
    }
}
],
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 9765032
| The FPolicy engine, FPolicy event or FPolicy policy specified already
exists

| 9765031
| If any of the FPolicy engine, FPolicy event, or FPolicy policy creation
fails due to a systematic error or hardware failure, the cause of the
failure is detailed in the error message

| 2621706
| The SVM UUID specified belongs to different SVM

| 2621462
| The SVM name specified does not exist
|===

```

== Definitions

[.api-def-first-level]

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#fpolicy_engine]
[.api-collapsible-fifth-title]
fpolicy_engine

The engine defines how ONTAP makes and manages connections to external
FPolicy servers.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Specifies the name to assign to the external server configuration.

|port
|integer
a|Port number of the FPolicy server application.
```

```
|primary_servers
|array[string]
a|
```

```
|secondary_servers
|array[string]
a|
```

```
|type
|string
a|The notification mode determines what ONTAP does after sending
notifications to FPolicy servers.
The possible values are:
```

```
***** synchronous - After sending a notification, wait for a response
from the FPolicy server.
```

```
***** asynchronous - After sending a notification, file request processing
continues.
```

```
* Default value: 1
* enum: ["synchronous", "asynchronous"]
* Introduced in: 9.6
```

```
|===
```

```
[#file_operations]
[.api-collapsible-fifth-title]
file_operations
```

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|close
|boolean
```

a|File close operations

|create

|boolean

a|File create operations

|create_dir

|boolean

a|Directory create operations

|delete

|boolean

a|File delete operations

|delete_dir

|boolean

a|Directory delete operations

|getattr

|boolean

a|Get attribute operations

|link

|boolean

a|Link operations

|lookup

|boolean

a|Lookup operations

|open

|boolean

a|File open operations

|read

|boolean

a|File read operations

```
|rename
|boolean
a|File rename operations
```

```
|rename_dir
|boolean
a|Directory rename operations
```

```
|setattr
|boolean
a|Set attribute operations
```

```
|symlink
|boolean
a|Symbolic link operations
```

```
|write
|boolean
a|File write operations
```

```
|===
```

```
[#filters]
[.api-collapsible-fifth-title]
filters
```

Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|close_with_modification
|boolean
```

a|Filter the client request for close with modification.

|close_with_read

|boolean

a|Filter the client request for close with read.

|close_without_modification

|boolean

a|Filter the client request for close without modification.

|exclude_directory

|boolean

a|Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.

|first_read

|boolean

a|Filter the client requests for the first-read.

|first_write

|boolean

a|Filter the client requests for the first-write.

|monitor_ads

|boolean

a|Filter the client request for alternate data stream.

|offline_bit

|boolean

a|Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.

|open_with_delete_intent

|boolean

a|Filter the client request for open with delete intent.

|open_with_write_intent

|boolean

a|Filter the client request for open with write intent.

|setattr_with_access_time_change

|boolean

a|Filter the client setattr requests for changing the access time of a file or directory.

|setattr_with_allocation_size_change

|boolean

a|Filter the client setattr requests for changing the allocation size of a file.

|setattr_with_creation_time_change

|boolean

a|Filter the client setattr requests for changing the creation time of a file or directory.

|setattr_with_dacl_change

|boolean

a|Filter the client setattr requests for changing dacl on a file or directory.

|setattr_with_group_change

|boolean

a|Filter the client setattr requests for changing group of a file or directory.

|setattr_with_mode_change

|boolean

a|Filter the client setattr requests for changing the mode bits on a file or directory.

|setattr_with_modify_time_change

|boolean

a|Filter the client setattr requests for changing the modification time of a file or directory.

|setattr_with_owner_change

|boolean

a|Filter the client setattr requests for changing owner of a file or directory.

|setattr_with_sacl_change

|boolean

a|Filter the client setattr requests for changing sacl on a file or directory.

|setattr_with_size_change

|boolean

a|Filter the client setattr requests for changing the size of a file.

|write_with_size_change

|boolean

a|Filter the client request for write with size change.

|===

[#fpolicy_event]

[.api-collapsible-fifth-title]

fpolicy_event

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|file_operations

|link:#file_operations[file_operations]

a|Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.

The event will check the operations specified from all client requests using the protocol.

|filters

|link:#filters[filters]

a|Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

|name

|string

a|Specifies the name of the FPolicy event.

|protocol

|string

a|Protocol for which event is created. If you specify protocol, then you must also specify a valid value for the file operation parameters.

The value of this parameter must be one of the following:

***** cifs - for the CIFS protocol.

***** nfsv3 - for the NFSv3 protocol.

***** nfsv4 - for the NFSv4 protocol.

|volume_monitoring

|boolean

a|Specifies whether volume operation monitoring is required.

|===

[#fpolicy_engine_reference]

[.api-collapsible-fifth-title]

fpolicy_engine_reference

FPolicy external engine

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

```
|string
a|The name of the FPolicy external engine.
```

```
|===
```

```
[#fpolicy_event_reference]
[.api-collapsible-fifth-title]
fpolicy_event_reference
```

FPolicy events

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|===
```

```
[#scope]
```

```
[.api-collapsible-fifth-title]
scope
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|exclude_export_policies
```

```
|array[string]
```

```
a|
```

```
|exclude_extension
```

```
|array[string]
```

```
a|
```

```
|exclude_shares
```

```
|array[string]
```

```

a|

|exclude_volumes
|array[string]
a|

|include_export_policies
|array[string]
a|

|include_extension
|array[string]
a|

|include_shares
|array[string]
a|

|include_volumes
|array[string]
a|

|===

[#fpolicy_policy]
[.api-collapsible-fifth-title]
fpolicy_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|engine
|link:#fpolicy_engine_reference[fpolicy_engine_reference]
a|FPolicy external engine

|events
|array[link:#fpolicy_event_reference[fpolicy_event_reference]]
a|

|mandatory
|boolean
a|Specifies what action to take on a file access event in a case when all

```

primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances.

|name
|string
a|Specifies the name of the policy.

|priority
|integer
a|Specifies the priority that is assigned to this policy.

|scope
|link:#scope[scope]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===

|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#fpolicy]
[.api-collapsible-fifth-title]

fpolicy

FPolicy is an infrastructure component of ONTAP that enables partner applications connected to your storage systems to monitor and set file access permissions. Every time a client accesses a file from a storage system, based on the configuration of FPolicy, the partner application is notified about file access.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|engines
|array[link:#fpolicy_engine[fpolicy_engine]]
a|

|events
|array[link:#fpolicy_event[fpolicy_event]]
a|

|policies
|array[link:#fpolicy_policy[fpolicy_policy]]
a|

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code
```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[IDd853e5ae3af882b0bfebe71d7001083f]]
= Delete the FPolicy configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}`#
```

Introduced In: 9.6

Deletes the FPolicy configuration for the specified SVM. Before deleting the FPolicy configuration, ensure that all policies belonging to the SVM are disabled.

== Related ONTAP commands

```
* `fpolicy delete`
* `fpolicy policy scope delete`
* `fpolicy policy delete`
* `fpolicy policy event delete`
* `fpolicy policy external-engine delete`
```

== Learn more

```
* xref:{relative_path}protocols_fpolicy_endpoint_overview.html[DOC
/protocols/fpolicy]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|===
```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 9765030

| Cannot delete an FPolicy configuration if any of the policy is enabled

| 9765031

| If any of the FPolicy engine, FPolicy event or FPolicy policy deletion fails due to a systemic error or hardware failure, the cause of the failure is detailed in the error message.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
```



```

    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

```

```

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

[[ID2e8bf8638e6ab540627dfb303b83a67b]]
= Retrieve the FPolicy configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}`#

*Introduced In:* 9.6

Retrieves an FPolicy configuration of an SVM.

== Related ONTAP commands

* `fpolicy show`
* `fpolicy policy show`
* `fpolicy policy scope show`
* `fpolicy policy event show`
* `fpolicy policy external-engine show`

```

== Learn more

* xref:{relative_path}protocols_fpolicy_endpoint_overview.html[DOC
/protocols/fpolicy]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

```

|engines
|array[link:#fpolicy_engine[fpolicy_engine]]
a|

|events
|array[link:#fpolicy_event[fpolicy_event]]
a|

|policies
|array[link:#fpolicy_policy[fpolicy_policy]]
a|

|svm
|link:#svm[svm]
a|

|===

```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "engines": [
    {
      "name": "fp_ex_eng",
      "port": "9876",
      "primary_servers": [
        "10.132.145.20",
        "10.140.101.109"
      ],
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "type": "string"
    }
  ],
  "events": [

```

```

    {
      "name": "event_nfs_close",
      "protocol": "string"
    }
  ],
  "policies": [
    {
      "engine": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string"
      },
      "events": [
        "event_nfs_close",
        "event_open"
      ],
      "name": "fp_policy_1",
      "scope": {
        "exclude_export_policies": [
          "string"
        ],
        "exclude_extension": [
          "string"
        ],
        "exclude_shares": [
          "string"
        ],
        "exclude_volumes": [
          "vol1",
          "vol_svm1",
          "*"
        ],
        "include_export_policies": [
          "string"
        ],
        "include_extension": [
          "string"
        ],
        "include_shares": [
          "sh1",
          "share_cifs"
        ],
        "include_volumes": [

```

```

        "vol1",
        "vol_svm1"
    ]
    }
}
],
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
    "error": {
        "arguments": [
            {
                "code": "string",
                "message": "string"
            }
        ]
    }
}

```

```

    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|self
|link:#href[href]
a|

```

```
|===
```

```
[#fpolicy_engine]  
[.api-collapsible-fifth-title]  
fpolicy_engine
```

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|Specifies the name to assign to the external server configuration.
```

```
|port  
|integer  
a|Port number of the FPolicy server application.
```

```
|primary_servers  
|array[string]  
a|
```

```
|secondary_servers  
|array[string]  
a|
```

```
|type  
|string  
a|The notification mode determines what ONTAP does after sending  
notifications to FPolicy servers.  
The possible values are:
```

```
***** synchronous - After sending a notification, wait for a response  
from the FPolicy server.
```

```
***** asynchronous - After sending a notification, file request processing  
continues.
```



```
* Default value: 1
* enum: ["synchronous", "asynchronous"]
* Introduced in: 9.6
```

```
|===
```

```
[#file_operations]
[.api-collapsible-fifth-title]
file_operations
```

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.

The event will check the operations specified from all client requests using the protocol.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|close
```

```
|boolean
```

```
a|File close operations
```

```
|create
```

```
|boolean
```

```
a|File create operations
```

```
|create_dir
```

```
|boolean
```

```
a|Directory create operations
```

```
|delete
```

```
|boolean
```

```
a|File delete operations
```

```
|delete_dir
```

```
|boolean
```

```
a|Directory delete operations
```

```
|getattr  
|boolean  
a|Get attribute operations
```

```
|link  
|boolean  
a|Link operations
```

```
|lookup  
|boolean  
a|Lookup operations
```

```
|open  
|boolean  
a|File open operations
```

```
|read  
|boolean  
a|File read operations
```

```
|rename  
|boolean  
a|File rename operations
```

```
|rename_dir  
|boolean  
a|Directory rename operations
```

```
|setattr  
|boolean  
a|Set attribute operations
```

```
|symlink  
|boolean  
a|Symbolic link operations
```

```
|write
|boolean
a|File write operations
```

```
|===
```

```
[#filters]
[.api-collapsible-fifth-title]
filters
```

Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|close_with_modification
|boolean
```

a|Filter the client request for close with modification.

```
|close_with_read
|boolean
```

a|Filter the client request for close with read.

```
|close_without_modification
|boolean
```

a|Filter the client request for close without modification.

```
|exclude_directory
|boolean
```

a|Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.

```
|first_read
|boolean
```

a|Filter the client requests for the first-read.

|first_write

|boolean

a|Filter the client requests for the first-write.

|monitor_ads

|boolean

a|Filter the client request for alternate data stream.

|offline_bit

|boolean

a|Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.

|open_with_delete_intent

|boolean

a|Filter the client request for open with delete intent.

|open_with_write_intent

|boolean

a|Filter the client request for open with write intent.

|setattr_with_access_time_change

|boolean

a|Filter the client setattr requests for changing the access time of a file or directory.

|setattr_with_allocation_size_change

|boolean

a|Filter the client setattr requests for changing the allocation size of a file.

|setattr_with_creation_time_change

|boolean

a|Filter the client setattr requests for changing the creation time of a file or directory.

```

|setattr_with_dacl_change
|boolean
a|Filter the client setattr requests for changing dacl on a file or
directory.

|setattr_with_group_change
|boolean
a|Filter the client setattr requests for changing group of a file or
directory.

|setattr_with_mode_change
|boolean
a|Filter the client setattr requests for changing the mode bits on a file
or directory.

|setattr_with_modify_time_change
|boolean
a|Filter the client setattr requests for changing the modification time of
a file or directory.

|setattr_with_owner_change
|boolean
a|Filter the client setattr requests for changing owner of a file or
directory.

|setattr_with_sacl_change
|boolean
a|Filter the client setattr requests for changing sacl on a file or
directory.

|setattr_with_size_change
|boolean
a|Filter the client setattr requests for changing the size of a file.

|write_with_size_change
|boolean
a|Filter the client request for write with size change.

|===

```

```
[#fpolicy_event]
[.api-collapsible-fifth-title]
fpolicy_event
```

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|file_operations
|link:#file_operations[file_operations]
a|Specifies the file operations for the FPolicy event. You must specify a
valid protocol in the protocol parameter.
The event will check the operations specified from all client requests
using the protocol.
```

```
|filters
|link:#filters[filters]
a|Specifies the list of filters for a given file operation for the
specified protocol.
When you specify the filters, you must specify the valid protocols and a
valid file operations.
```

```
|name
|string
a|Specifies the name of the FPolicy event.
```

```
|protocol
|string
a|Protocol for which event is created. If you specify protocol, then you
must also specify a valid value for the file operation parameters.
The value of this parameter must be one of the following:
```

```
***** cifs - for the CIFS protocol.
```

```
***** nfsv3 - for the NFSv3 protocol.
```

***** nfsv4 - for the NFSv4 protocol.

|volume_monitoring

|boolean

a|Specifies whether volume operation monitoring is required.

|===

[#fpolicy_engine_reference]

[.api-collapsible-fifth-title]

fpolicy_engine_reference

FPolicy external engine

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the FPolicy external engine.

|===

[#fpolicy_event_reference]

[.api-collapsible-fifth-title]

fpolicy_event_reference

FPolicy events

[cols=3*,options=header]

|===

|Name

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|===

[#scope]
[.api-collapsible-fifth-title]
scope

[cols=3*,options=header]
|===
|Name
|Type
|Description

|exclude_export_policies
|array[string]
a|

|exclude_extension
|array[string]
a|

|exclude_shares
|array[string]
a|

|exclude_volumes
|array[string]
a|

|include_export_policies
|array[string]
a|

|include_extension
|array[string]
a|

```



```

|include_shares
|array[string]
a|

|include_volumes
|array[string]
a|

|===

[#fpolicy_policy]
[.api-collapsible-fifth-title]
fpolicy_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies if the policy is enabled on the SVM or not. If no value is
mentioned for this field but priority is set, then this policy will be
enabled.

|engine
|link:#fpolicy_engine_reference[fpolicy_engine_reference]
a|FPolicy external engine

|events
|array[link:#fpolicy_event_reference[fpolicy_event_reference]]
a|

|mandatory
|boolean
a|Specifies what action to take on a file access event in a case when all
primary and secondary servers are down or no response is received from the
FPolicy servers within a given timeout period. When this parameter is set
to true, file access events will be denied under these circumstances.

|name

```

```

|string
a|Specifies the name of the policy.

|priority
|integer
a|Specifies the priority that is assigned to this policy.

|scope
|link:#scope[scope]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage FPolicy engine configuration
```

```
:leveloffset: +1
```

```
[[ID1bc9653dd2e42f33b71b0ec784e0b2a5]]
```

```
= Protocols fpolicy svm.uuid engines endpoint overview
```

```
== Overview
```

The FPolicy engine allows you to configure the external servers to which the file access notifications are sent. As part of FPolicy engine configuration, you can configure the server(s) to which the notification is sent, an optional set of secondary server(s) to which the notification is sent in the case of the primary server(s) failure, the port number for FPolicy application and the type of the engine, synchronous or asynchronous.

For the synchronous engine, ONTAP will wait for a response from the FPolicy application before it allows the operation. With an asynchronous engine, ONTAP proceeds with the operation processing after sending the notification to the FPolicy application. An engine can belong to multiple FPolicy policies.

```
== Examples
```

```
=== Creating an FPolicy engine
```

```
'''
```

```
----
```

```

# The API:
POST /protocols/fpolicy/{svm.uuid}/engines

#The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/engines/" -H "accept: application/json" -H "Content-Type: application/json" -d '{"name": "engine0", "port": 9876, "primary_servers": [ "10.132.145.22", "10.140.101.109" ], "secondary_servers": [ "10.132.145.20", "10.132.145.21" ], "type": "synchronous"}'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "name": "engine0",
      "primary_servers": [
        "10.132.145.22",
        "10.140.101.109"
      ],
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "port": 9876,
      "type": "synchronous"
    }
  ]
}
-----

'''

=== Creating an FPolicy engine with the minimum required fields

'''

-----

# The API:
POST /protocols/fpolicy/{svm.uuid}/engines

#The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/engines/" -H "accept: application/json" -H "Content-

```

```
Type: application/json" -d "{ \"name\": \"engine0\", \"port\": 9876,
\"primary_servers\": [ \"10.132.145.22\", \"10.140.101.109\" ], \"type\":
\"synchronous\"}"
```

The response:

```
{
  "num_records": 1,
  "records": [
    {
      "name": "engine0",
      "primary_servers": [
        "10.132.145.22",
        "10.140.101.109"
      ],
      "port": 9876,
      "type": "synchronous"
    }
  ]
}
```

'''

=== Retrieving an FPolicy engine configuration for a particular SVM

'''

The API:

GET /protocols/fpolicy/{svm.uuid}/engines

The call:

```
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/engines/?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"
```

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
      },
      "name": "cifs",
      "primary_servers": [
```

```

        "10.20.20.10"
    ],
    "port": 9876,
    "type": "synchronous"
},
{
    "svm": {
        "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
    },
    "name": "nfs",
    "primary_servers": [
        "10.23.140.64",
        "10.140.101.109"
    ],
    "secondary_servers": [
        "10.132.145.20",
        "10.132.145.22"
    ],
    "port": 9876,
    "type": "synchronous"
}
],
"num_records": 2
}
----

'''

=== Retrieving a specific FPolicy engine configuration for an SVM

'''

----

# The Api:
GET /protocols/fpolicy/{svm.uuid}/engines/{name}

#The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/engines/cifs?fields=*" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
  },

```

```

"name": "cifs",
"primary_servers": [
    "10.20.20.10"
],
"port": 9876,
"type": "synchronous"
}
----

'''

=== Updating an FPolicy engine for an SVM

'''

----

# The API:
PATCH /protocols/fpolicy/{svm.uuid}/engines/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/engines/cifs" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"port\": 6666, \"secondary_servers\": [ \"10.132.145.20\", \"10.132.145.21\" ], \"type\": \"synchronous\"}"
----

'''

=== Updating all the attributes of a specific FPolicy engine for an SVM

'''

----

# The API:
PATCH /protocols/fpolicy/{svm.uuid}/engines/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/engines/cifs" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"port\": 9876, \"primary_servers\": [ \"10.132.145.20\", \"10.140.101.109\" ], \"secondary_servers\": [ \"10.132.145.23\", \"10.132.145.21\" ], \"type\": \"synchronous\"}"
----

'''

```


=== Deleting a specific FPolicy engine for an SVM

'''

The API:

DELETE /protocols/fpolicy/{svm.uuid}/engines/{name}

The call:

curl -X DELETE "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/events/cifs" -H "accept: application/json"

'''

[[IDcbff08f2585d637a0bf9ca93e5a802ee]]

= Retrieve the FPolicy engine configuration for all engines of an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/fpolicy/{svm.uuid}/engines`#

Introduced In: 9.6

Retrieves FPolicy engine configurations of all the engines for a specified SVM. ONTAP allows creation of cluster-level FPolicy engines that act as a template for all the SVMs belonging to the cluster. These cluster-level FPolicy engines are also retrieved for the specified SVM.

== Related ONTAP commands

* `fpolicy policy external-engine show`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_engines_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/engines]

== Parameters

```

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|secondary_servers
|string
|query
|False
a|Filter by secondary_servers

|port
|integer
|query
|False
a|Filter by port

|type
|string
|query
|False
a|Filter by type

|name
|string
|query
|False
a|Filter by name

|primary_servers
|string
|query
|False
a|Filter by primary_servers

|svm.uuid
|string
|path

```

```

|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.

```

Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of records

|records

|array[link:#fpolicy_engine[fpolicy_engine]]

a|

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "_links": {

 "next": {

 "href": "/api/resourcelink"

 },

 "self": {

 "href": "/api/resourcelink"

 }

 },

 "records": [

 {

 "name": "fp_ex_eng",

```

    "port": "9876",
    "primary_servers": [
      "10.132.145.20",
      "10.140.101.109"
    ],
    "secondary_servers": [
      "10.132.145.20",
      "10.132.145.21"
    ],
    "type": "string"
  }
]
}
====

== Error

```

Status: Default,

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",

```

```

    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

```

```
|===
```

```
[#fpolicy_engine]  
[.api-collapsible-fifth-title]  
fpolicy_engine
```

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string
```

a|Specifies the name to assign to the external server configuration.

```
|port  
|integer
```

a|Port number of the FPolicy server application.

```
|primary_servers  
|array[string]
```

a|

```
|secondary_servers  
|array[string]
```

a|

```
|type  
|string
```

a|The notification mode determines what ONTAP does after sending notifications to FPolicy servers.

The possible values are:

***** synchronous - After sending a notification, wait for a response from the FPolicy server.

***** asynchronous - After sending a notification, file request processing continues.

- * Default value: 1
- * enum: ["synchronous", "asynchronous"]
- * Introduced in: 9.6

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID82ee760834b11c42cca8abeef6c6d203]]
= Create an FPolicy engine configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}/engines`#

*Introduced In:* 9.6

Creates an FPolicy engine configuration for a specified SVM. FPolicy
engine creation is allowed only on data SVMs.

== Required properties

* `svm.uuid` - Existing SVM in which to create the FPolicy engine.
* `name` - Name of external engine.
* `port` - Port number of the FPolicy server application.
* `primary_servers` - List of primary FPolicy servers to which the node
will send notifications.

== Recommended optional properties

* `secondary_servers` - It is recommended to configure secondary FPolicy
server to which the node will send notifications when the primary server
is down.

```

== Default property values

* `type` - `_synchronous_`

== Related ONTAP commands

* ``fpolicy policy external-engine create``

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_engines_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/engines]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|===

== Request Body

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Specifies the name to assign to the external server configuration.

|port
|integer
a|Port number of the FPolicy server application.

|primary_servers
|array[string]
a|

|secondary_servers
|array[string]
a|

|type
|string
a|The notification mode determines what ONTAP does after sending
notifications to FPolicy servers.
The possible values are:

***** synchronous - After sending a notification, wait for a response
from the FPolicy server.

***** asynchronous - After sending a notification, file request processing
continues.

* Default value: 1
* enum: ["synchronous", "asynchronous"]
* Introduced in: 9.6

|===

.Example request
[%collapsible%closed]

```

```
====  
[source,json,subs+=macros]  
{  
  "name": "fp_ex_eng",  
  "port": "9876",  
  "primary_servers": [  
    "10.132.145.20",  
    "10.140.101.109"  
  ],  
  "secondary_servers": [  
    "10.132.145.20",  
    "10.132.145.21"  
  ],  
  "type": "string"  
}  
====  
  
== Response
```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#fpolicy_engine[fpolicy_engine]]
a|

|===


.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "name": "fp_ex_eng",
      "port": "9876",
      "primary_servers": [
        "10.132.145.20",
        "10.140.101.109"
      ],
      "secondary_servers": [
        "10.132.145.20",
        "10.132.145.21"
      ],
      "type": "string"
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 9764885

| The primary secondary server has a redundant IP address

| 9764953

| The name of the FPolicy engine is "native" which is reserved by the system

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#fpolicy_engine]

[.api-collapsible-fifth-title]

fpolicy_engine

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Specifies the name to assign to the external server configuration.

|port

|integer

a|Port number of the FPolicy server application.

|primary_servers

|array[string]

```

a|

|secondary_servers
|array[string]
a|

|type
|string
a|The notification mode determines what ONTAP does after sending
notifications to FPolicy servers.
The possible values are:

***** synchronous - After sending a notification, wait for a response
from the FPolicy server.

***** asynchronous - After sending a notification, file request processing
continues.

* Default value: 1
* enum: ["synchronous", "asynchronous"]
* Introduced in: 9.6

|===

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#error_arguments]

```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message


```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID2f0e432746728bce14d7c49fa376173f]]
= Delete an FPolicy external engine configuration
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}/engines/{name}`#
```

Introduced In: 9.6

Deletes the FPolicy external engine configuration. Deletion of an FPolicy engine that is attached to one or more FPolicy policies is not allowed.

== Related ONTAP commands

* `fpolicy policy external-engine modify`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_engines_endpoint_overview.h
tml[DOC /protocols/fpolicy/{svm.uuid}/engines]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|name
```

```

|string
|path
|True
a|

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 9764940
| At least one FPolicy policy is using the FPolicy engine

| 9764887
| The FPolicy engine is a cluster level FPolicy engine
|===

```

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string

```

```

a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDf2d6bf84153a3bd4d435800986597d73]]
= Retrieve a particular FPolicy engine configuration for an SVM

```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/fpolicy/{svm.uuid}/engines/{name}`#
```

Introduced In: 9.6

Retrieves a particular FPolicy engine configuration of a specified SVM. A cluster-level FPolicy engine configuration cannot be retrieved for a data SVM.

== Related ONTAP commands

* `fpolicy policy external-engine show`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_engines_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/engines]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|name
|string
|path
|True
a|
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
```

```
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Specifies the name to assign to the external server configuration.

|port
|integer
a|Port number of the FPolicy server application.

|primary_servers
|array[string]
a|

|secondary_servers
|array[string]
a|

|type
|string
a|The notification mode determines what ONTAP does after sending
notifications to FPolicy servers.
The possible values are:

***** synchronous - After sending a notification, wait for a response
from the FPolicy server.

***** asynchronous - After sending a notification, file request processing
continues.

* Default value: 1
```

```
* enum: ["synchronous", "asynchronous"]
```

```
* Introduced in: 9.6
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{  
  "name": "fp_ex_eng",  
  "port": "9876",  
  "primary_servers": [  
    "10.132.145.20",  
    "10.140.101.109"  
  ],  
  "secondary_servers": [  
    "10.132.145.20",  
    "10.132.145.21"  
  ],  
  "type": "string"  
}
```

```
====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
====
```

```

[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]

```



```

[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID77ad22f87bdaa2303813336ac2979b36]]
= Update a specific FPolicy engine configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}/engines/{name}`#

*Introduced In:* 9.6

```

Updates a specific FPolicy engine configuration of an SVM. Modification of an FPolicy engine that is attached to one or more enabled FPolicy policies

is not allowed.

== Related ONTAP commands

* `fpolicy policy external-engine modify`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_engines_endpoint_overview.h
tml[DOC /protocols/fpolicy/{svm.uuid}/engines]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|name

|string

|path

|True

a|

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

```

|port
|integer
a|Port number of the FPolicy server application.

|primary_servers
|array[string]
a|

|secondary_servers
|array[string]
a|

|type
|string
a|The notification mode determines what ONTAP does after sending
notifications to FPolicy servers.
The possible values are:

***** synchronous - After sending a notification, wait for a response
from the FPolicy server.

***** asynchronous - After sending a notification, file request processing
continues.

* Default value: 1
* enum: ["synchronous", "asynchronous"]
* Introduced in: 9.6

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "port": "9876",
  "primary_servers": [
    "10.132.145.20",
    "10.140.101.109"
  ],
  "secondary_servers": [
    "10.132.145.20",
    "10.132.145.21"

```

```

    },
    "type": "string"
}
====

```

```

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 9764922
| The primary and secondary server has a redundant IP address
|
| 9764942
| At least one FPolicy policy is using the FPolicy engine
|
| 9764886
| FPolicy engine is a cluster-level FPolicy engine
|====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#fpolicy_engine]
[.api-collapsible-fifth-title]
fpolicy_engine

```

The engine defines how ONTAP makes and manages connections to external FPolicy servers.

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|port
|integer
a|Port number of the FPolicy server application.

|primary_servers
|array[string]
a|

|secondary_servers
|array[string]
a|

|type
|string
a|The notification mode determines what ONTAP does after sending
notifications to FPolicy servers.
The possible values are:

***** synchronous - After sending a notification, wait for a response
from the FPolicy server.

***** asynchronous - After sending a notification, file request processing
continues.

* Default value: 1
* enum: ["synchronous", "asynchronous"]
* Introduced in: 9.6

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage FPolicy event configuration
```

```
:leveloffset: +1
```

```
[[IDff7d87a588d03e012ff3621fad9d9a48]]
```

```
= Protocols fpolicy svm.uuid events endpoint overview
```

```
== Overview
```

FPolicy events configurations allow you to specify which file access is monitored. As part of an FPolicy event, you can configure the SVM for which the events are generated, the name of the event configuration, the protocol (cifs, nfsv3/nfsv4) for which the events are generated, the file operations which are monitored, and filters that can be used to filter the unwanted notification generation for a specified protocol and file operation.

Each protocol has a set of supported file operations and filters. An SVM can have multiple events. A single FPolicy policy can have multiple FPolicy events.

```
== Examples
```

```
=== Creating an FPolicy event for a CIFS protocol with all the supported  
file operations and filters
```

```
'''
```

```
----
```

```
# The API:
```

```
POST /protocols/fpolicy/{svm.uuid}/events
```

```
# The call:
```

```
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-  
ae49-0050568e2c1e/events?return_records=true" -H "accept:
```

```
application/json" -H "Content-Type: application/json" -d "{
  \"file_operations\": { \"close\": true, \"create\": true, \"create_dir\":
true, \"delete\": true, \"delete_dir\": true, \"getattr\": true, \"open\":
true, \"read\": true, \"rename\": true, \"rename_dir\": true, \"setattr\":
true, \"write\": true }, \"filters\": { \"close_with_modification\": true,
\"close_with_read\": true, \"close_without_modification\": true,
\"first_read\": true, \"first_write\": true, \"monitor_ads\": true,
\"offline_bit\": true, \"open_with_delete_intent\": true,
\"open_with_write_intent\": true, \"write_with_size_change\": true },
\"name\": \"event_cifs\", \"protocol\": \"cifs\", \"volume_monitoring\":
true}"
```

The response:

```
{
  "num_records": 1,
  "records": [
    {
      "name": "event_cifs",
      "protocol": "cifs",
      "volume_monitoring": true,
      "file_operations": {
        "close": true,
        "create": true,
        "create_dir": true,
        "delete": true,
        "delete_dir": true,
        "getattr": true,
        "open": true,
        "read": true,
        "write": true,
        "rename": true,
        "rename_dir": true,
        "setattr": true
      },
      "filters": {
        "monitor_ads": true,
        "close_with_modification": true,
        "close_without_modification": true,
        "close_with_read": true,
        "first_read": true,
        "first_write": true,
        "offline_bit": true,
        "open_with_delete_intent": true,
        "open_with_write_intent": true,
        "write_with_size_change": true
      }
    }
  ]
}
```



```

    }
  ]
}
----

'''

=== Creating an FPolicy event for an NFS protocol with all the supported
file operations and filters

'''

----

# The API:
post /protocols/fpolicy/{svm.uuid}/events

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-
ae49-0050568e2c1e/events?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{
  \"file_operations\": { \"create\": true, \"create_dir\": true, \"delete\":
true, \"delete_dir\": true, \"link\": true, \"lookup\": true, \"read\":
true, \"rename\": true, \"rename_dir\": true, \"setattr\": true,
  \"symlink\": true, \"write\": true }, \"filters\": { \"offline_bit\":
true, \"write_with_size_change\": true }, \"name\": \"event_nfsv3\",
  \"protocol\": \"nfsv3\", \"volume_monitoring\": false}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "name": "event_nfsv3",
      "protocol": "nfsv3",
      "volume_monitoring": false,
      "file_operations": {
        "create": true,
        "create_dir": true,
        "delete": true,
        "delete_dir": true,
        "link": true,
        "lookup": true,
        "read": true,
        "write": true,
        "rename": true,

```

```

        "rename_dir": true,
        "setattr": true,
        "symlink": true
    },
    "filters": {
        "offline_bit": true,
        "write_with_size_change": true
    }
}
]
}

```

'''

=== Retrieving all of the FPolicy event configurations for a specified SVM

'''

The API:

GET /protocols/fpolicy/{svm.uuid}/events

The call:

```

curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/events/?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"

```

The response:

```

{
  "records": [
    {
      "svm": {
        "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
      },
      "name": "cluster",
      "protocol": "cifs",
      "volume_monitoring": false,
      "file_operations": {
        "close": true,
        "create": false,
        "create_dir": false,
        "delete": false,
        "delete_dir": false,
        "setattr": false,

```

```

    "link": false,
    "lookup": false,
    "open": false,
    "read": false,
    "write": false,
    "rename": false,
    "rename_dir": false,
    "setattr": false,
    "symlink": false
  },
  "filters": {
    "monitor_ads": false,
    "close_with_modification": false,
    "close_without_modification": false,
    "close_with_read": true,
    "first_read": false,
    "first_write": false,
    "offline_bit": false,
    "open_with_delete_intent": false,
    "open_with_write_intent": false,
    "write_with_size_change": false,
    "setattr_with_owner_change": false,
    "setattr_with_group_change": false,
    "setattr_with_sacl_change": false,
    "setattr_with_dacl_change": false,
    "setattr_with_modify_time_change": false,
    "setattr_with_access_time_change": false,
    "setattr_with_creation_time_change": false,
    "setattr_with_mode_change": false,
    "setattr_with_size_change": false,
    "setattr_with_allocation_size_change": false,
    "exclude_directory": false
  }
},
{
  "svm": {
    "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
  },
  "name": "event_cifs",
  "protocol": "cifs",
  "volume_monitoring": true,
  "file_operations": {
    "close": true,
    "create": true,
    "create_dir": true,
    "delete": true,

```

```

    "delete_dir": true,
    "getattr": true,
    "link": false,
    "lookup": false,
    "open": true,
    "read": true,
    "write": true,
    "rename": true,
    "rename_dir": true,
    "setattr": true,
    "symlink": false
},
"filters": {
    "monitor_ads": true,
    "close_with_modification": true,
    "close_without_modification": true,
    "close_with_read": true,
    "first_read": true,
    "first_write": true,
    "offline_bit": true,
    "open_with_delete_intent": true,
    "open_with_write_intent": true,
    "write_with_size_change": true,
    "setattr_with_owner_change": false,
    "setattr_with_group_change": false,
    "setattr_with_sacl_change": false,
    "setattr_with_dacl_change": false,
    "setattr_with_modify_time_change": false,
    "setattr_with_access_time_change": false,
    "setattr_with_creation_time_change": false,
    "setattr_with_mode_change": false,
    "setattr_with_size_change": false,
    "setattr_with_allocation_size_change": false,
    "exclude_directory": false
}
},
"num_records": 2
}
----

'''

=== Retrieving a specific FPolicy event configuration for an SVM

'''

```

The API:

GET /protocols/fpolicy/{svm.uuid}/events/{name}

The call:

```
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/events/event_cifs?fields=*&return_records=true&return_timeout=15" -H "accept: application/json"
```

The response:

```
{
  "svm": {
    "uuid": "4f643fb4-fd21-11e8-ae49-0050568e2c1e"
  },
  "name": "event_cifs",
  "protocol": "cifs",
  "volume_monitoring": true,
  "file_operations": {
    "close": true,
    "create": true,
    "create_dir": true,
    "delete": true,
    "delete_dir": true,
    "getattr": true,
    "link": false,
    "lookup": false,
    "open": true,
    "read": true,
    "write": true,
    "rename": true,
    "rename_dir": true,
    "setattr": true,
    "symlink": false
  },
  "filters": {
    "monitor_ads": true,
    "close_with_modification": true,
    "close_without_modification": true,
    "close_with_read": true,
    "first_read": true,
    "first_write": true,
    "offline_bit": true,
    "open_with_delete_intent": true,
  }
}
```

```

"open_with_write_intent": true,
"write_with_size_change": true,
"setattr_with_owner_change": false,
"setattr_with_group_change": false,
"setattr_with_sacl_change": false,
"setattr_with_dacl_change": false,
"setattr_with_modify_time_change": false,
"setattr_with_access_time_change": false,
"setattr_with_creation_time_change": false,
"setattr_with_mode_change": false,
"setattr_with_size_change": false,
"setattr_with_allocation_size_change": false,
"exclude_directory": false
}
}
----

'''

=== Updating a specific FPolicy event configuration for a specified SVM

'''

----

# The API:
PATCH /protocols/fpolicy/{svm.uuid}/events/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/events/event_cifs" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"file_operations\": { \"close\": false, \"create\": false, \"read\": true }, \"filters\": { \"close_with_modification\": false, \"close_with_read\": false, \"close_without_modification\": false }, \"protocol\": \"cifs\", \"volume_monitoring\": false}"
----

'''

=== Deleting a specific FPolicy event configuration for a specific SVM

'''

----

# The API:

```

```
DELETE /protocols/fpolicy/{svm.uuid}/events/{name}
```

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/protocols/fpolicy/4f643fb4-fd21-11e8-ae49-0050568e2c1e/events/event_cifs" -H "accept: application/json"
----
```

'''

```
[[IDe87c64c338189bf577ad9fe8798a75ca]]
```

= Retrieve the FPolicy event configuration for all events for an SVM

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/fpolicy/{svm.uuid}/events`#
```

Introduced In: 9.6

Retrieves FPolicy event configurations for all events for a specified SVM. ONTAP allows the creation of cluster-level FPolicy events that act as a template for all the data SVMs belonging to the cluster. These cluster-level FPolicy events are also retrieved for the specified SVM.

== Related ONTAP commands

* `fpolicy policy event show`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_events_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/events]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|protocol
|string
|query
|False
a|Filter by protocol
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|filters.close_with_modification
|boolean
|query
|False
a|Filter by filters.close_with_modification
```

```
|filters.close_without_modification
|boolean
|query
|False
a|Filter by filters.close_without_modification
```

```
|filters.offline_bit
|boolean
|query
|False
a|Filter by filters.offline_bit
```

```
|filters.write_with_size_change
|boolean
|query
|False
a|Filter by filters.write_with_size_change
```

```
|filters.setattr_with_creation_time_change
|boolean
|query
|False
```



```
a|Filter by filters.setattr_with_creation_time_change
```

```
|filters.setattr_with_size_change
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Filter by filters.setattr_with_size_change
```

```
|filters.open_with_delete_intent
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Filter by filters.open_with_delete_intent
```

```
|filters.setattr_with_owner_change
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Filter by filters.setattr_with_owner_change
```

```
|filters.setattr_with_access_time_change
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Filter by filters.setattr_with_access_time_change
```

```
|filters.setattr_with_group_change
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Filter by filters.setattr_with_group_change
```

```
|filters.setattr_with_dacl_change
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Filter by filters.setattr_with_dacl_change
```

```
|filters.setattr_with_allocation_size_change
```

```
|boolean
```

```
|query
|False
a|Filter by filters.setattr_with_allocation_size_change
```

```
|filters.exclude_directory
|boolean
|query
|False
a|Filter by filters.exclude_directory
```

```
|filters.open_with_write_intent
|boolean
|query
|False
a|Filter by filters.open_with_write_intent
```

```
|filters.monitor_ads
|boolean
|query
|False
a|Filter by filters.monitor_ads
```

```
|filters.setattr_with_mode_change
|boolean
|query
|False
a|Filter by filters.setattr_with_mode_change
```

```
|filters.setattr_with_sacl_change
|boolean
|query
|False
a|Filter by filters.setattr_with_sacl_change
```

```
|filters.setattr_with_modify_time_change
|boolean
|query
|False
a|Filter by filters.setattr_with_modify_time_change
```

```
|filters.first_read
|boolean
|query
|False
a|Filter by filters.first_read
```

```
|filters.close_with_read
|boolean
|query
|False
a|Filter by filters.close_with_read
```

```
|filters.first_write
|boolean
|query
|False
a|Filter by filters.first_write
```

```
|volume_monitoring
|boolean
|query
|False
a|Filter by volume_monitoring
```

```
|file_operations.close
|boolean
|query
|False
a|Filter by file_operations.close
```

```
|file_operations.lookup
|boolean
|query
|False
a|Filter by file_operations.lookup
```

```
|file_operations.delete_dir
|boolean
|query
|False
a|Filter by file_operations.delete_dir
```

```
|file_operations.create_dir
|boolean
|query
|False
a|Filter by file_operations.create_dir
```

```
|file_operations.create
|boolean
|query
|False
a|Filter by file_operations.create
```

```
|file_operations.setattr
|boolean
|query
|False
a|Filter by file_operations.setattr
```

```
|file_operations.getattr
|boolean
|query
|False
a|Filter by file_operations.getattr
```

```
|file_operations.delete
|boolean
|query
|False
a|Filter by file_operations.delete
```

```
|file_operations.symlink
|boolean
|query
|False
a|Filter by file_operations.symlink
```

```
|file_operations.rename
|boolean
|query
```

```

|False
a|Filter by file_operations.rename

|file_operations.open
|boolean
|query
|False
a|Filter by file_operations.open

|file_operations.read
|boolean
|query
|False
a|Filter by file_operations.read

|file_operations.rename_dir
|boolean
|query
|False
a|Filter by file_operations.rename_dir

|file_operations.link
|boolean
|query
|False
a|Filter by file_operations.link

|file_operations.write
|boolean
|query
|False
a|Filter by file_operations.write

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields

```

```

|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```



```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#fpolicy_event[fpolicy_event]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "name": "event_nfs_close",
      "protocol": "string"
    }
  ]
}
====

== Error

```



```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]

```

```
|===  
|Name  
|Type  
|Description
```

```
|href  
|string  
a|
```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links
```

```
[cols=3*,options=header]
```

```
|===  
|Name  
|Type  
|Description
```

```
|next  
|link:#href[href]  
a|
```

```
|self  
|link:#href[href]  
a|
```

```
|===
```

```
[#file_operations]  
[.api-collapsible-fifth-title]  
file_operations
```

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

```
[cols=3*,options=header]
```

```
|===  
|Name  
|Type
```

```
|Description

|close
|boolean
a|File close operations

|create
|boolean
a|File create operations

|create_dir
|boolean
a|Directory create operations

|delete
|boolean
a|File delete operations

|delete_dir
|boolean
a|Directory delete operations

|getattr
|boolean
a|Get attribute operations

|link
|boolean
a|Link operations

|lookup
|boolean
a|Lookup operations

|open
|boolean
a|File open operations
```

```
|read
|boolean
a|File read operations
```

```
|rename
|boolean
a|File rename operations
```

```
|rename_dir
|boolean
a|Directory rename operations
```

```
|setattr
|boolean
a|Set attribute operations
```

```
|symlink
|boolean
a|Symbolic link operations
```

```
|write
|boolean
a|File write operations
```

```
|===
```

```
[#filters]
[.api-collapsible-fifth-title]
filters
```

Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

|Description

|close_with_modification

|boolean

a|Filter the client request for close with modification.

|close_with_read

|boolean

a|Filter the client request for close with read.

|close_without_modification

|boolean

a|Filter the client request for close without modification.

|exclude_directory

|boolean

a|Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.

|first_read

|boolean

a|Filter the client requests for the first-read.

|first_write

|boolean

a|Filter the client requests for the first-write.

|monitor_ads

|boolean

a|Filter the client request for alternate data stream.

|offline_bit

|boolean

a|Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.

|open_with_delete_intent

|boolean

a|Filter the client request for open with delete intent.

|open_with_write_intent

|boolean

a|Filter the client request for open with write intent.

|setattr_with_access_time_change

|boolean

a|Filter the client setattr requests for changing the access time of a file or directory.

|setattr_with_allocation_size_change

|boolean

a|Filter the client setattr requests for changing the allocation size of a file.

|setattr_with_creation_time_change

|boolean

a|Filter the client setattr requests for changing the creation time of a file or directory.

|setattr_with_dacl_change

|boolean

a|Filter the client setattr requests for changing dacl on a file or directory.

|setattr_with_group_change

|boolean

a|Filter the client setattr requests for changing group of a file or directory.

|setattr_with_mode_change

|boolean

a|Filter the client setattr requests for changing the mode bits on a file or directory.

|setattr_with_modify_time_change

|boolean

a|Filter the client setattr requests for changing the modification time of a file or directory.

```
|setattr_with_owner_change
|boolean
a|Filter the client setattr requests for changing owner of a file or
directory.
```

```
|setattr_with_sacl_change
|boolean
a|Filter the client setattr requests for changing sacl on a file or
directory.
```

```
|setattr_with_size_change
|boolean
a|Filter the client setattr requests for changing the size of a file.
```

```
|write_with_size_change
|boolean
a|Filter the client request for write with size change.
```

```
|===
```

```
[#fpolicy_event]
[.api-collapsible-fifth-title]
fpolicy_event
```

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|file_operations
|link:#file_operations[file_operations]
a|Specifies the file operations for the FPolicy event. You must specify a
valid protocol in the protocol parameter.
The event will check the operations specified from all client requests
```

using the protocol.

|filters

|link:#filters[filters]

a|Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

|name

|string

a|Specifies the name of the FPolicy event.

|protocol

|string

a|Protocol for which event is created. If you specify protocol, then you must also specify a valid value for the file operation parameters.

The value of this parameter must be one of the following:

***** cifs - for the CIFS protocol.

***** nfsv3 - for the NFSv3 protocol.

***** nfsv4 - for the NFSv4 protocol.

|volume_monitoring

|boolean

a|Specifies whether volume operation monitoring is required.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description


```

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

```

=====

[[IDeb059fb05eelba162f62192480c4ecab]]

= Create the FPolicy event configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/protocols/fpolicy/{svm.uuid}/events`#

Introduced In: 9.6

Creates an FPolicy event configuration for a specified SVM. FPolicy event creation is allowed only on data SVMs. When a protocol is specified, you must specify a file operation or a file operation and filters.

== Required properties

- * `svm.uuid` - Existing SVM in which to create the FPolicy event.
- * `name` - Name of the FPolicy event.

== Recommended optional properties

- * `file-operations` - List of file operations to monitor.
- * `protocol` - Protocol for which the file operations should be monitored.
- * `filters` - List of filters for the specified file operations.

== Default property values

If not specified in POST, the following default property values are assigned:

- * `file_operations.+++` - `_false_`
- * `filters.+++` - `_false_`
- * `volume-monitoring` - `_false_`

== Related ONTAP commands

- * `fpolicy policy event create`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_events_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/events]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False

a|The default is false. If set to true, the records are returned.

* Default value:

|svm.uuid
|string
|path
|True

a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===

|Name
|Type
|Description

|file_operations
|link:#file_operations[file_operations]

a|Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

|filters

```

|link:#filters[filters]
a|Specifies the list of filters for a given file operation for the
specified protocol.
When you specify the filters, you must specify the valid protocols and a
valid file operations.

|name
|string
a|Specifies the name of the FPolicy event.

|protocol
|string
a|Protocol for which event is created. If you specify protocol, then you
must also specify a valid value for the file operation parameters.
The value of this parameter must be one of the following:

***** cifs - for the CIFS protocol.

***** nfsv3 - for the NFSv3 protocol.

***** nfsv4 - for the NFSv4 protocol.

|volume_monitoring
|boolean
a|Specifies whether volume operation monitoring is required.

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "name": "event_nfs_close",
  "protocol": "string"
}
=====

== Response

```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#fpolicy_event[fpolicy_event]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "name": "event_nfs_close",
      "protocol": "string"
    }
  ]
}
====

== Error

```

Status: Default

```

|===
| Error Code | Description

| 9764929
| The file operation is not supported by the protocol

| 9764955
| The filter is not supported by the protocol

```

```
| 9764930
| The filter is not supported by any of the file operations

| 9764946
| The protocol is specified without a file operation or a file operation
and filter pair
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#file_operations]
[.api-collapsible-fifth-title]
file_operations
```

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.

The event will check the operations specified from all client requests using the protocol.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|close
|boolean
a|File close operations
```

```
|create
|boolean
a|File create operations
```

```
|create_dir
|boolean
a|Directory create operations
```

|delete
|boolean
a|File delete operations

|delete_dir
|boolean
a|Directory delete operations

|getattr
|boolean
a|Get attribute operations

|link
|boolean
a|Link operations

|lookup
|boolean
a|Lookup operations

|open
|boolean
a|File open operations

|read
|boolean
a|File read operations

|rename
|boolean
a|File rename operations

|rename_dir
|boolean
a|Directory rename operations

|setattr
|boolean

a|Set attribute operations

|symlink

|boolean

a|Symbolic link operations

|write

|boolean

a|File write operations

|===

[#filters]

[.api-collapsible-fifth-title]

filters

Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|close_with_modification

|boolean

a|Filter the client request for close with modification.

|close_with_read

|boolean

a|Filter the client request for close with read.

|close_without_modification

|boolean

a|Filter the client request for close without modification.


```

|exclude_directory
|boolean
a|Filter the client requests for directory operations. When this filter is
specified directory operations are not monitored.

|first_read
|boolean
a|Filter the client requests for the first-read.

|first_write
|boolean
a|Filter the client requests for the first-write.

|monitor_ads
|boolean
a|Filter the client request for alternate data stream.

|offline_bit
|boolean
a|Filter the client request for offline bit set. FPolicy server receives
notification only when offline files are accessed.

|open_with_delete_intent
|boolean
a|Filter the client request for open with delete intent.

|open_with_write_intent
|boolean
a|Filter the client request for open with write intent.

|setattr_with_access_time_change
|boolean
a|Filter the client setattr requests for changing the access time of a
file or directory.

|setattr_with_allocation_size_change
|boolean
a|Filter the client setattr requests for changing the allocation size of a
file.

```

|setattr_with_creation_time_change

|boolean

a|Filter the client setattr requests for changing the creation time of a file or directory.

|setattr_with_dacl_change

|boolean

a|Filter the client setattr requests for changing dacl on a file or directory.

|setattr_with_group_change

|boolean

a|Filter the client setattr requests for changing group of a file or directory.

|setattr_with_mode_change

|boolean

a|Filter the client setattr requests for changing the mode bits on a file or directory.

|setattr_with_modify_time_change

|boolean

a|Filter the client setattr requests for changing the modification time of a file or directory.

|setattr_with_owner_change

|boolean

a|Filter the client setattr requests for changing owner of a file or directory.

|setattr_with_sacl_change

|boolean

a|Filter the client setattr requests for changing sacl on a file or directory.

|setattr_with_size_change

|boolean

a|Filter the client setattr requests for changing the size of a file.

```
|write_with_size_change
|boolean
a|Filter the client request for write with size change.
```

```
|===
```

```
[#fpolicy_event]
[.api-collapsible-fifth-title]
fpolicy_event
```

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|file_operations
|link:#file_operations[file_operations]
a|Specifies the file operations for the FPolicy event. You must specify a
valid protocol in the protocol parameter.
The event will check the operations specified from all client requests
using the protocol.
```

```
|filters
|link:#filters[filters]
a|Specifies the list of filters for a given file operation for the
specified protocol.
When you specify the filters, you must specify the valid protocols and a
valid file operations.
```

```
|name
|string
a|Specifies the name of the FPolicy event.
```

```
|protocol
```

|string
a|Protocol for which event is created. If you specify protocol, then you must also specify a valid value for the file operation parameters. The value of this parameter must be one of the following:

***** cifs - for the CIFS protocol.

***** nfsv3 - for the NFSv3 protocol.

***** nfsv4 - for the NFSv4 protocol.

|volume_monitoring
|boolean
a|Specifies whether volume operation monitoring is required.

|===

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name

```

|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

```

```
//end collapsible .Definitions block
====
```

```
[[IDc5207cd341b4086aded70a989a14794b]]
```

```
= Delete a specific FPolicy event configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/protocols/fpolicy/{svm.uuid}/events/{name}`#
```

```
*Introduced In:* 9.6
```

Deletes a specific FPolicy event configuration for an SVM. A cluster-level FPolicy event configuration cannot be modified for a data SVM through REST. An FPolicy event that is attached to an FPolicy policy cannot be deleted.

```
== Related ONTAP commands
```

```
* `fpolicy policy event delete`
```

```
== Learn more
```

```
*
```

```
xref:{relative_path}protocols_fpolicy_svm.uuid_events_endpoint_overview.ht  
ml[DOC /protocols/fpolicy/{svm.uuid}/events]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|name
```

```
|string
```

```
|path
```

```
|True
```

```
a|
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
|===
| Error Code | Description

| 9764874
| The FPolicy event is a cluster event

| 9764947
| The FPolicy event is attached to an FPolicy policy
|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
```

```

[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]

```



```
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDca1d563ec4553f29a810acdb2a65a1c3]]
= Retrieve a specific FPolicy event configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}/events/{name}`#

*Introduced In:* 9.6

Retrieves a specific FPolicy event configuration for an SVM. A cluster-
level FPolicy event configuration cannot be retrieved for a data SVM
```

through a REST API.

== Related ONTAP commands

* `fpolicy policy event show`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_events_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/events]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|name

|string

|path

|True

a|

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|file_operations
|link:#file_operations[file_operations]
a|Specifies the file operations for the FPolicy event. You must specify a
valid protocol in the protocol parameter.
The event will check the operations specified from all client requests
using the protocol.

|filters
|link:#filters[filters]
a|Specifies the list of filters for a given file operation for the
specified protocol.
When you specify the filters, you must specify the valid protocols and a
valid file operations.

|name
|string
a|Specifies the name of the FPolicy event.

|protocol
|string
a|Protocol for which event is created. If you specify protocol, then you
must also specify a valid value for the file operation parameters.
The value of this parameter must be one of the following:

***** cifs - for the CIFS protocol.

***** nfsv3 - for the NFSv3 protocol.

***** nfsv4 - for the NFSv4 protocol.

|volume_monitoring
|boolean
a|Specifies whether volume operation monitoring is required.
```

```
|===
```

.Example response

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{  
  "name": "event_nfs_close",  
  "protocol": "string"  
}
```

```
=====
```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

.Example error

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{  
  "error": {  
    "arguments": [  
      {  
        "code": "string",  
        "message": "string"  
      }  
    ],  
    "code": "4",  
    "message": "entry doesn't exist",  
    "target": "uuid"  
  }  
}
```

```

}
}
====

```

```

== Definitions

```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#file_operations]
[.api-collapsible-fifth-title]
file_operations

```

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.

The event will check the operations specified from all client requests using the protocol.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|close
|boolean
a|File close operations

```

```

|create
|boolean
a|File create operations

```

```

|create_dir
|boolean
a|Directory create operations

```

```

|delete
|boolean
a|File delete operations

```

|delete_dir
|boolean
a|Directory delete operations

|getattr
|boolean
a|Get attribute operations

|link
|boolean
a|Link operations

|lookup
|boolean
a|Lookup operations

|open
|boolean
a|File open operations

|read
|boolean
a|File read operations

|rename
|boolean
a|File rename operations

|rename_dir
|boolean
a|Directory rename operations

|setattr
|boolean
a|Set attribute operations

|symlink
|boolean

a|Symbolic link operations

|write

|boolean

a|File write operations

|===

[#filters]

[.api-collapsible-fifth-title]

filters

Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|close_with_modification

|boolean

a|Filter the client request for close with modification.

|close_with_read

|boolean

a|Filter the client request for close with read.

|close_without_modification

|boolean

a|Filter the client request for close without modification.

|exclude_directory

|boolean

a|Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.

|first_read
|boolean
a|Filter the client requests for the first-read.

|first_write
|boolean
a|Filter the client requests for the first-write.

|monitor_ads
|boolean
a|Filter the client request for alternate data stream.

|offline_bit
|boolean
a|Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.

|open_with_delete_intent
|boolean
a|Filter the client request for open with delete intent.

|open_with_write_intent
|boolean
a|Filter the client request for open with write intent.

|setattr_with_access_time_change
|boolean
a|Filter the client setattr requests for changing the access time of a file or directory.

|setattr_with_allocation_size_change
|boolean
a|Filter the client setattr requests for changing the allocation size of a file.

|setattr_with_creation_time_change
|boolean
a|Filter the client setattr requests for changing the creation time of a

file or directory.

|setattr_with_dacl_change

|boolean

a|Filter the client setattr requests for changing dacl on a file or directory.

|setattr_with_group_change

|boolean

a|Filter the client setattr requests for changing group of a file or directory.

|setattr_with_mode_change

|boolean

a|Filter the client setattr requests for changing the mode bits on a file or directory.

|setattr_with_modify_time_change

|boolean

a|Filter the client setattr requests for changing the modification time of a file or directory.

|setattr_with_owner_change

|boolean

a|Filter the client setattr requests for changing owner of a file or directory.

|setattr_with_sacl_change

|boolean

a|Filter the client setattr requests for changing sacl on a file or directory.

|setattr_with_size_change

|boolean

a|Filter the client setattr requests for changing the size of a file.

|write_with_size_change

|boolean

a|Filter the client request for write with size change.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[IDf7075fd56591ffc777183bf52fda44f4]]
= Update a specific FPolicy event configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}/events/{name}`#
```

***Introduced In:* 9.6**

Updates a specific FPolicy event configuration for an SVM. A cluster-level FPolicy event configuration cannot be modified for a data SVM through REST. When the file operations and filters fields are modified, the previous values are retained and new values are added to the list of previous values. To remove a particular file operation or filter, set its value to false in the request.

== Related ONTAP commands

* `fpolicy policy event modify`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_events_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/events]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|name
```

```
|string
```

```
|path
```

```
|True
```

```
a|
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the SVM to which this object belongs.
```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|file_operations
```

```
|link:#file_operations[file_operations]
```

```
a|Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
```

```
The event will check the operations specified from all client requests using the protocol.
```

```
|filters
```

```
|link:#filters[filters]
```

```
a|Specifies the list of filters for a given file operation for the specified protocol.
```

```
When you specify the filters, you must specify the valid protocols and a valid file operations.
```

```
|protocol
|string
a|Protocol for which event is created. If you specify protocol, then you
must also specify a valid value for the file operation parameters.
The value of this parameter must be one of the following:
```

```
***** cifs - for the CIFS protocol.
```

```
***** nfsv3 - for the NFSv3 protocol.
```

```
***** nfsv4 - for the NFSv4 protocol.
```

```
|volume_monitoring
|boolean
a|Specifies whether volume operation monitoring is required.
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "protocol": "string"
}
=====
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
|===
| Error Code | Description

| 9764873
| The event is a cluster event
```

```

| 9764929
| The file operation is not supported by the protocol

| 9764955
| The filter is not supported by the protocol

| 9764930
| The filter is not supported by any of the file operations

| 9764946
| The protocol is specifed without file operation or a file operation and
filter pair
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#file_operations]
[.api-collapsible-fifth-title]
file_operations

```

Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.
The event will check the operations specified from all client requests using the protocol.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|close
|boolean
a|File close operations

|create
|boolean
a|File create operations

```

|create_dir
|boolean
a|Directory create operations

|delete
|boolean
a|File delete operations

|delete_dir
|boolean
a|Directory delete operations

|getattr
|boolean
a|Get attribute operations

|link
|boolean
a|Link operations

|lookup
|boolean
a|Lookup operations

|open
|boolean
a|File open operations

|read
|boolean
a|File read operations

|rename
|boolean
a|File rename operations

```
|rename_dir
|boolean
a|Directory rename operations
```

```
|setattr
|boolean
a|Set attribute operations
```

```
|symlink
|boolean
a|Symbolic link operations
```

```
|write
|boolean
a|File write operations
```

```
|===
```

```
[#filters]
[.api-collapsible-fifth-title]
filters
```

Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|close_with_modification
|boolean
a|Filter the client request for close with modification.
```

```
|close_with_read
|boolean
a|Filter the client request for close with read.
```


|close_without_modification

|boolean

a|Filter the client request for close without modification.

|exclude_directory

|boolean

a|Filter the client requests for directory operations. When this filter is specified directory operations are not monitored.

|first_read

|boolean

a|Filter the client requests for the first-read.

|first_write

|boolean

a|Filter the client requests for the first-write.

|monitor_ads

|boolean

a|Filter the client request for alternate data stream.

|offline_bit

|boolean

a|Filter the client request for offline bit set. FPolicy server receives notification only when offline files are accessed.

|open_with_delete_intent

|boolean

a|Filter the client request for open with delete intent.

|open_with_write_intent

|boolean

a|Filter the client request for open with write intent.

|setattr_with_access_time_change

|boolean

a|Filter the client setattr requests for changing the access time of a

file or directory.

|setattr_with_allocation_size_change

|boolean

a|Filter the client setattr requests for changing the allocation size of a file.

|setattr_with_creation_time_change

|boolean

a|Filter the client setattr requests for changing the creation time of a file or directory.

|setattr_with_dacl_change

|boolean

a|Filter the client setattr requests for changing dacl on a file or directory.

|setattr_with_group_change

|boolean

a|Filter the client setattr requests for changing group of a file or directory.

|setattr_with_mode_change

|boolean

a|Filter the client setattr requests for changing the mode bits on a file or directory.

|setattr_with_modify_time_change

|boolean

a|Filter the client setattr requests for changing the modification time of a file or directory.

|setattr_with_owner_change

|boolean

a|Filter the client setattr requests for changing owner of a file or directory.

|setattr_with_sacl_change

|boolean

a|Filter the client setattr requests for changing sacl on a file or directory.

|setattr_with_size_change

|boolean

a|Filter the client setattr requests for changing the size of a file.

|write_with_size_change

|boolean

a|Filter the client request for write with size change.

|===

[#fpolicy_event]

[.api-collapsible-fifth-title]

fpolicy_event

The information that a FPolicy process needs to determine what file access operations to monitor and for which of the monitored events notifications should be sent to the external FPolicy server.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|file_operations

|link:#file_operations[file_operations]

a|Specifies the file operations for the FPolicy event. You must specify a valid protocol in the protocol parameter.

The event will check the operations specified from all client requests using the protocol.

|filters

|link:#filters[filters]

a|Specifies the list of filters for a given file operation for the specified protocol.

When you specify the filters, you must specify the valid protocols and a valid file operations.

|protocol
|string
a|Protocol for which event is created. If you specify protocol, then you must also specify a valid value for the file operation parameters. The value of this parameter must be one of the following:

***** cifs - for the CIFS protocol.

***** nfsv3 - for the NFSv3 protocol.

***** nfsv4 - for the NFSv4 protocol.

|volume_monitoring
|boolean
a|Specifies whether volume operation monitoring is required.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]

```

error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====

:leveloffset: -1

= Manage SVM FPolicy configuration

:leveloffset: +1

[[ID3cda65ae221d59edd6743b83aecc0195]]
= Protocols fpolicy svm.uuid policies endpoint overview

```

== Overview

The FPolicy policy acts as a container for different constituents of the FPolicy such as FPolicy events and the FPolicy engine. It also provides a platform for policy management functions, such as policy enabling and disabling. As part of FPolicy policy configuration, you can specify the name of policy, the SVM to which it belongs, the FPolicy events to monitor, the FPolicy engine to which the generated notifications are sent and the policy priority. FPolicy policy configuration also allows you to configure the file access behaviour when the primary and secondary servers are down. Under such circumstances, if the "mandatory" field is set to true, file access is denied.

Each FPolicy policy is associated with a scope which allows you to restrain the scope of the policy to specified storage objects such as volume, shares and export or to a set of file extensions such as .txt, .jpeg. An FPolicy policy can be configured to send notifications, to the FPolicy server or for native file blocking which uses the file extension specified in the policy scope. An SVM can have multiple FPolicy policies which can be enabled or disabled independently of each other.

== Examples

=== Creating an FPolicy policy

Use the following API to create an FPolicy policy configuration. Note that the `_return_records=true_` query parameter used to obtain the newly created entry in the response.

```
'''
```

```
----
```

```
# The API:
```

```
POST /protocols/fpolicy/{svm.uuid}/policies
```

```
# The call:
```

```
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-b64a-0050568eeb34/policies?return_records=true" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"engine\": { \"name\": \"engine1\" }, \"events\": [ { \"name\": \"cifs\" }, { \"name\": \"nfs\" } ], \"mandatory\": true, \"name\": \"FPolicy_policy_0\", \"scope\": { \"exclude_export_policies\": [ \"export_poll\" ], \"exclude_extension\": [ \"txt\", \"png\" ], \"exclude_shares\": [ \"sh1\" ], \"exclude_volumes\": [ \"vol0\" ], \"include_export_policies\": [
```

```

\"export_pol10\" ], \"include_extension\": [ \"pdf\" ],
\"include_shares\": [ \"sh2\", \"sh3\" ], \"include_volumes\": [ \"vol1\",
\"vol2\" ] } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "name": "FPolicy_policy_0",
      "events": [
        {
          "name": "cifs"
        },
        {
          "name": "nfs"
        }
      ],
      "engine": {
        "name": "engine1"
      },
      "scope": {
        "include_shares": [
          "sh2",
          "sh3"
        ],
        "exclude_shares": [
          "sh1"
        ],
        "include_volumes": [
          "vol1",
          "vol2"
        ],
        "exclude_volumes": [
          "vol0"
        ],
        "include_export_policies": [
          "export_pol10"
        ],
        "exclude_export_policies": [
          "export_pol1"
        ],
        "include_extension": [
          "pdf"
        ],
        "exclude_extension": [

```

```

        "txt",
        "png"
    ]
},
"mandatory": true
}
]
}
----

'''

=== Creating and enable an FPolicy policy

'''

----

# The API:
POST /protocols/fpolicy/{svm.uuid}/policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"priority\":
1, \"engine\": { \"name\": \"engine1\" }, \"events\": [ { \"name\":
\"cifs\" }, { \"name\": \"nfs\" } ], \"mandatory\": true, \"name\":
\"FPolicy_policy_on\", \"scope\": { \"exclude_export_policies\": [
\"export_pol1\" ], \"exclude_extension\": [ \"txt\", \"png\" ],
\"exclude_shares\": [ \"sh1\" ], \"exclude_volumes\": [ \"vol0\" ],
\"include_export_policies\": [ \"export_pol10\" ], \"include_extension\":
[ \"pdf\" ], \"include_shares\": [ \"sh2\", \"sh3\" ],
\"include_volumes\": [ \"vol1\", \"vol2\" ] } }"

# The response:
{
"num_records": 1,
"records": [
{
"name": "FPolicy_policy_0",
"priority": 1,
"events": [
{
"name": "cifs"
},
{

```



```

        "name": "nfs"
    }
],
"engine": {
    "name": "engine1"
},
"scope": {
    "include_shares": [
        "sh2",
        "sh3"
    ],
    "exclude_shares": [
        "sh1"
    ],
    "include_volumes": [
        "vol1",
        "vol2"
    ],
    "exclude_volumes": [
        "vol0"
    ],
    "include_export_policies": [
        "export_pol10"
    ],
    "exclude_export_policies": [
        "export_pol1"
    ],
    "include_extension": [
        "pdf"
    ],
    "exclude_extension": [
        "txt",
        "png"
    ]
},
"mandatory": true
}
]
}
----

'''

```

=== Creating an FPolicy policy with the minimum required fields and a native engine

```
'''
----

# The API:
POST /protocols/fpolicy/{svm.uuid}/policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-
b64a-0050568eeb34/policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"events\": [
{ \"name\": \"cifs\" }, { \"name\": \"nfs\" } ], \"name\":
\"pol_minimum_fields\", \"scope\": { \"include_volumes\": [ \"vol1\",
\"vol2\" ] } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "name": "pol_minimum_fields",
      "events": [
        {
          "name": "cifs"
        },
        {
          "name": "nfs"
        }
      ],
      "scope": {
        "include_volumes": [
          "vol1",
          "vol2"
        ]
      }
    }
  ]
}
'''

=== Retrieving all the FPolicy policy configurations for an SVM

'''
```

```

-----

# The API:
GET /protocols/fpolicy/{svm.uuid}/policies

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-b64a-0050568eeb34/policies?fields=*&return_records=true&return_timeout=15"
-H "accept: application/json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
      },
      "name": "pol0",
      "enabled": false,
      "events": [
        {
          "name": "cifs"
        },
        {
          "name": "nfs"
        }
      ],
      "engine": {
        "name": "engine1"
      },
      "scope": {
        "include_shares": [
          "sh2",
          "sh3"
        ],
        "exclude_shares": [
          "sh1"
        ],
        "include_volumes": [
          "vol1",
          "vol2"
        ],
        "exclude_volumes": [
          "vol0"
        ],
        "include_export_policies": [

```

```

        "export_policy10"
    ],
    "exclude_export_policies": [
        "export_policy1"
    ],
    "include_extension": [
        "pdf"
    ],
    "exclude_extension": [
        "txt",
        "png"
    ]
},
"mandatory": true
},
{
    "svm": {
        "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "FPolicy_policy_on",
    "enabled": true,
    "priority": 1,
    "events": [
        {
            "name": "cifs"
        },
        {
            "name": "nfs"
        }
    ],
    "engine": {
        "name": "engine1"
    },
    "scope": {
        "include_shares": [
            "sh2",
            "sh3"
        ],
        "exclude_shares": [
            "sh1"
        ],
        "include_volumes": [
            "vol1",
            "vol2"
        ],
        "exclude_volumes": [

```

```

        "vol0"
    ],
    "include_export_policies": [
        "export_pol10"
    ],
    "exclude_export_policies": [
        "export_pol1"
    ],
    "include_extension": [
        "pdf"
    ],
    "exclude_extension": [
        "txt",
        "png"
    ]
},
"mandatory": true
},
{
    "svm": {
        "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "cluster_pol",
    "enabled": false,
    "events": [
        {
            "name": "cluster"
        }
    ],
    "engine": {
        "name": "native"
    },
    "mandatory": true
},
{
    "svm": {
        "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
    },
    "name": "pol_minimum_fields",
    "enabled": false,
    "events": [
        {
            "name": "cifs"
        },
        {
            "name": "nfs"
        }
    ]
}

```

```

    }
  ],
  "engine": {
    "name": "native"
  },
  "scope": {
    "include_volumes": [
      "vol1",
      "vol2"
    ]
  },
  "mandatory": true
}
],
"num_records": 4
}
----
'''

```

=== Retrieving all of the FPolicy policy configurations for the FPolicy engine "engine1" for an SVM

```
'''
```

```
----
```

The API:

```
GET /protocols/fpolicy/{svm.uuid}/policies/{name}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-b64a-0050568eeb34/policies?engine.name=engine1&fields=*&return_records=true&return_timeout=15" -H "accept: application/json"
```

The response:

```

{
  "records": [
    {
      "svm": {
        "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
      },
      "name": "pol0",
      "enabled": false,
      "events": [

```

```

    {
      "name": "cifs"
    },
    {
      "name": "nfs"
    }
  ],
  "engine": {
    "name": "engine1"
  },
  "scope": {
    "include_export_policies": [
      "export_poll0"
    ],
    "exclude_export_policies": [
      "export_poll1"
    ],
    "include_extension": [
      "pdf"
    ],
    "exclude_extension": [
      "txt",
      "png"
    ]
  },
  "mandatory": true
},
{
  "svm": {
    "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
  },
  "name": "FPolicy_policy_on",
  "enabled": true,
  "priority": 1,
  "events": [
    {
      "name": "cifs"
    },
    {
      "name": "nfs"
    }
  ],
  "engine": {
    "name": "engine1"
  },
  "scope": {

```

```

    "include_shares": [
        "sh2",
        "sh3"
    ],
    "exclude_shares": [
        "sh1"
    ],
    "include_volumes": [
        "vol1",
        "vol2"
    ],
    "exclude_volumes": [
        "vol0"
    ],
    "include_export_policies": [
        "export_pol10"
    ],
    "exclude_export_policies": [
        "export_pol1"
    ],
    "include_extension": [
        "pdf"
    ],
    "exclude_extension": [
        "txt",
        "png"
    ]
},
"mandatory": true
}
],
"num_records": 2
}
----

'''

=== Retrieving a particular FPolicy policy configuration for an SVM

'''

----

# The API:
GET /protocols/fpolicy/{svm.uuid}/policies/{name}

```



```
# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-b64a-0050568eeb34/policies/pol0" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "a00fac5d-0164-11e9-b64a-0050568eeb34"
  },
  "name": "pol0",
  "enabled": false,
  "events": [
    {
      "name": "cifs"
    },
    {
      "name": "nfs"
    }
  ],
  "engine": {
    "name": "engine1"
  },
  "scope": {
    "include_shares": [
      "sh2",
      "sh3"
    ],
    "exclude_shares": [
      "sh1"
    ],
    "include_volumes": [
      "vol1",
      "vol2"
    ],
    "exclude_volumes": [
      "vol0"
    ],
    "include_export_policies": [
      "export_pol10"
    ],
    "exclude_export_policies": [
      "export_pol1"
    ],
    "include_extension": [
      "pdf"
    ],
  ],
}
```

```

    "exclude_extension": [
        "txt",
        "png"
    ]
},
"mandatory": true
}
----

'''

=== Updating a particular FPolicy policy

'''

----

# The API:
PATCH /protocols/fpolicy/{svm.uuid}/policies/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-b64a-0050568eeb34/policies/pol0" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"engine\": { \"name\": \"native\"
}, \"events\": [ { \"name\": \"cifs\" } ], \"mandatory\": false,
\"scope\": { \"include_volumes\": [ \"*\" ] } }"
----

'''

=== Enabling a particular FPolicy policy

'''

----

# The API:
PATCH /protocols/fpolicy/{svm.uuid}/policies/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-b64a-0050568eeb34/policies/pol0" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"enabled\": true, \"priority\":
3}"
----

'''

```

=== Disabling a particular FPolicy policy

'''

The API:

PATCH /protocols/fpolicy/{svm.uuid}/policies/{name}

The call:

curl -X PATCH "https://<mgmt-ip>/api/protocols/fpolicy/a00fac5d-0164-11e9-b64a-0050568eeb34/policies/pol0" -H "accept: application/json" -H "Content-Type: application/json" -d '{"enabled": true }'

'''

[[IDdb0dc323091ad46fa4457d9d3b63b1be]]

= Retrieve the FPolicy configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/fpolicy/{svm.uuid}/policies`#

Introduced In: 9.6

Retrieves the FPolicy policy configuration of an SVM. ONTAP allows the creation of a cluster level FPolicy policy that acts as a template for all the data SVMs belonging to the cluster. This cluster level FPolicy policy is also retrieved for the specified SVM.

== Related ONTAP commands

* `fpolicy policy show`

* `fpolicy policy scope show`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_policies_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/policies]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|priority
|integer
|query
|False
a|Filter by priority

|engine.name
|string
|query
|False
a|Filter by engine.name

|events.name
|string
|query
|False
a|Filter by events.name

|scope.include_shares
|string
|query
|False
a|Filter by scope.include_shares

|scope.include_extension
|string
|query
|False
a|Filter by scope.include_extension

|scope.include_volumes

```
|string
|query
|False
a|Filter by scope.include_volumes

|scope.include_export_policies
|string
|query
|False
a|Filter by scope.include_export_policies

|scope.exclude_export_policies
|string
|query
|False
a|Filter by scope.exclude_export_policies

|scope.exclude_shares
|string
|query
|False
a|Filter by scope.exclude_shares

|scope.exclude_volumes
|string
|query
|False
a|Filter by scope.exclude_volumes

|scope.exclude_extension
|string
|query
|False
a|Filter by scope.exclude_extension

|name
|string
|query
|False
a|Filter by name
```

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|mandatory
|boolean
|query
|False
a|Filter by mandatory
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

```
* Default value: 1
```

```
|return_timeout
```

```

|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of Records

|records
|array[link:#fpolicy_policy[fpolicy_policy]]
a|

|===

```

```
.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "engine": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "string"
      },
      "events": [
        "event_nfs_close",
        "event_open"
      ],
      "name": "fp_policy_1",
      "scope": {
        "exclude_export_policies": [
          "string"
        ],
        "exclude_extension": [
          "string"
        ],
        "exclude_shares": [
          "string"
        ],
        "exclude_volumes": [
          "vol1",
          "vol_svm1",
          "*"
        ],
        "include_export_policies": [

```



```

        "string"
      ],
      "include_extension": [
        "string"
      ],
      "include_shares": [
        "sh1",
        "share_cifs"
      ],
      "include_volumes": [
        "vol1",
        "vol_svm1"
      ]
    }
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",

```

```

        "message": "string"
    }
],
"code": "4",
"message": "entry doesn't exist",
"target": "uuid"
}
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]

```

```

a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#fpolicy_engine_reference]
[.api-collapsible-fifth-title]
fpolicy_engine_reference

FPolicy external engine

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the FPolicy external engine.

```

```
|===
```

```
[#fpolicy_event_reference]  
[.api-collapsible-fifth-title]  
fpolicy_event_reference
```

FPolicy events

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|===
```

```
[#scope]
```

```
[.api-collapsible-fifth-title]  
scope
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|exclude_export_policies
```

```
|array[string]
```

```
a|
```

```
|exclude_extension
```

```
|array[string]
```

```
a|
```

```
|exclude_shares
```

```

|array[string]
a|

|exclude_volumes
|array[string]
a|

|include_export_policies
|array[string]
a|

|include_extension
|array[string]
a|

|include_shares
|array[string]
a|

|include_volumes
|array[string]
a|

|===

[#fpolicy_policy]
[.api-collapsible-fifth-title]
fpolicy_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies if the policy is enabled on the SVM or not. If no value is
mentioned for this field but priority is set, then this policy will be
enabled.

|engine
|link:#fpolicy_engine_reference[fpolicy_engine_reference]
a|FPolicy external engine

```

```

|events
|array[link:#fpolicy_event_reference[fpolicy_event_reference]]
a|

|mandatory
|boolean
a|Specifies what action to take on a file access event in a case when all
primary and secondary servers are down or no response is received from the
FPolicy servers within a given timeout period. When this parameter is set
to true, file access events will be denied under these circumstances.

|name
|string
a|Specifies the name of the policy.

|priority
|integer
a|Specifies the priority that is assigned to this policy.

|scope
|link:#scope[scope]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string

```

```

a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID0c863bf80989bba0dd3297ca91fb7820]]
= Create the FPolicy configuration for an SVM

```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/protocols/fpolicy/{svm.uuid}/policies`#
```

Introduced In: 9.6

Creates an FPolicy policy configuration for the specified SVM. To create an FPolicy policy, you must specify the policy scope and the FPolicy events to be monitored.

Important notes:

- * A single policy can monitor multiple events.
- * An FPolicy engine is an optional field whose default value is set to native. A native engine can be used to simply block the file access based on the file extensions specified in the policy scope.
- * To enable a policy, the policy priority must be specified. If the priority is not specified, the policy is created but it is not enabled.
- * The "mandatory" field, if set to true, blocks the file access when the primary or secondary FPolicy servers are down.

== Required properties

- * `svm.uuid` - Existing SVM in which to create the FPolicy policy.
- * `events` - Name of the events to monitor.
- * `name` - Name of the FPolicy policy.
- * `scope` - Scope of the policy. Can be limited to exports, volumes, shares or file extensions.
- * `priority` - Priority of the policy (ranging from 1 to 10).

== Default property values

- * `mandatory` - `_true_`
- * `engine` - `_native_`

== Related ONTAP commands

- * ``fpolicy policy scope create``
- * ``fpolicy policy create``
- * ``fpolicy enable``

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_policies_endpoint_overview.
html[DOC /protocols/fpolicy/{svm.uuid}/policies]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False

a|The default is false. If set to true, the records are returned.

* Default value:

|svm.uuid
|string
|path
|True

a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===

|Name
|Type
|Description

|engine
|link:#fpolicy_engine_reference[fpolicy_engine_reference]

a|FPolicy external engine

|events
|array[link:#fpolicy_event_reference[fpolicy_event_reference]]
a|

|mandatory
|boolean
a|Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances.

|name
|string
a|Specifies the name of the policy.

|priority
|integer
a|Specifies the priority that is assigned to this policy.

|scope
|link:#scope[scope]
a|

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
 "enabled": null,
 "engine": {
 "name": "string"
 },
 "events": [
 "event_nfs_close",
 "event_open"
],
 "name": "fp_policy_1",
 "scope": {
 "exclude_export_policies": [
 "string"
],
 "exclude_extension": [
 "string"
],
 },
}

```

    "exclude_shares": [
        "string"
    ],
    "exclude_volumes": [
        "vol1",
        "vol_svm1",
        "*"
    ],
    "include_export_policies": [
        "string"
    ],
    "include_extension": [
        "string"
    ],
    "include_shares": [
        "sh1",
        "share_cifs"
    ],
    "include_volumes": [
        "vol1",
        "vol_svm1"
    ]
}
}
====

```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of Records

|records
|array[link:#fpolicy_policy[fpolicy_policy]]
a|

```

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "records": [
    {
      "enabled": null,
      "engine": {
        "name": "string"
      },
      "events": [
        "event_nfs_close",
        "event_open"
      ],
      "name": "fp_policy_1",
      "scope": {
        "exclude_export_policies": [
          "string"
        ],
        "exclude_extension": [
          "string"
        ],
        "exclude_shares": [
          "string"
        ],
        "exclude_volumes": [
          "vol1",
          "vol_svm1",
          "*"
        ],
        "include_export_policies": [
          "string"
        ],
        "include_extension": [
          "string"
        ],
        "include_shares": [
          "sh1",
          "share_cifs"
        ],
        "include_volumes": [
          "vol1",
```

```

        "vol_svm1"
    ]
}
}
]
}
====

```

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 9765027
| FPolicy creation is successful but it cannot be enabled as the priority
is already in use by another policy
|
| 9764898
| An FPolicy policy cannot be created without defining its scope
|====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|====
|Name
|Type
|Description

|href
|string
a|

```

|===

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#fpolicy_engine_reference]
[.api-collapsible-fifth-title]
fpolicy_engine_reference
```

FPolicy external engine

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the FPolicy external engine.
```

|===

```
[#fpolicy_event_reference]
[.api-collapsible-fifth-title]
fpolicy_event_reference
```

FPolicy events

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|name
|string
a|
```

|===

```
[#scope]
[.api-collapsible-fifth-title]
scope
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|exclude_export_policies
|array[string]
a|
```

```
|exclude_extension
|array[string]
a|
```

```
|exclude_shares
|array[string]
a|
```

```
|exclude_volumes
|array[string]
a|
```

```
|include_export_policies
|array[string]
a|
```

```
|include_extension
|array[string]
a|
```

```
|include_shares
|array[string]
a|
```

```
|include_volumes
|array[string]
a|
```

```
|===
```

```
[#fpolicy_policy]
```

```

[.api-collapsible-fifth-title]
fpolicy_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|engine
|link:#fpolicy_engine_reference[fpolicy_engine_reference]
a|FPolicy external engine

|events
|array[link:#fpolicy_event_reference[fpolicy_event_reference]]
a|

|mandatory
|boolean
a|Specifies what action to take on a file access event in a case when all
primary and secondary servers are down or no response is received from the
FPolicy servers within a given timeout period. When this parameter is set
to true, file access events will be denied under these circumstances.

|name
|string
a|Specifies the name of the policy.

|priority
|integer
a|Specifies the priority that is assigned to this policy.

|scope
|link:#scope[scope]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```



```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

|===

//end collapsible .Definitions block

====

[[ID068d01356d0f50c1dfd41dc3398f9af2]]

= Delete an FPolicy configuration for an SVM

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/protocols/fpolicy/{svm.uuid}/policies/{name}`#

Introduced In: 9.6

Deletes a particular FPolicy policy configuration for a specified SVM. To delete a policy, you must first disable the policy.

== Related ONTAP commands

* `fpolicy policy scope delete`

* `fpolicy policy delete`

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_policies_endpoint_overview.html[DOC /protocols/fpolicy/{svm.uuid}/policies]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|name

|string

|path

```

|True
a|

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

```

ONTAP Error Response Codes

|===
| Error Code | Description

| 9764900
| Deletion of a cluster level FPolicy policy is not supported

| 9764941
| Cannot delete an enabled FPolicy policy
|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID73a4613893b8d3752598a1c64b062d04]]
= Retrieve an FPolicy configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}/policies/{name}`#

```

Introduced In: 9.6

Retrieves a particular FPolicy policy configuration for a specified SVM. Cluster-level FPolicy policy configuration details cannot be retrieved for a data SVM.

== Related ONTAP commands

- * ``fpolicy policy show``
- * ``fpolicy policy scope show``
- * ``fpolicy show``

== Learn more

*

xref:{relative_path}protocols_fpolicy_svm.uuid_policies_endpoint_overview.
html[DOC /protocols/fpolicy/{svm.uuid}/policies]

== Parameters

[cols=5*,options=header]

|===

Name
Type
In
Required
Description

name
string
path
True
a

svm.uuid
string
path
True
a UUID of the SVM to which this object belongs.

fields
array[string]
query
False

```
a|Specify the fields to return.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|Specifies if the policy is enabled on the SVM or not. If no value is mentioned for this field but priority is set, then this policy will be enabled.
```

```
|engine
```

```
|link:#fpolicy_engine_reference[fpolicy_engine_reference]
```

```
a|FPolicy external engine
```

```
|events
```

```
|array[link:#fpolicy_event_reference[fpolicy_event_reference]]
```

```
a|
```

```
|mandatory
```

```
|boolean
```

```
a|Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances.
```

```
|name
```

```
|string
```

```
a|Specifies the name of the policy.
```

```
|priority
```

```
|integer
```

```
a|Specifies the priority that is assigned to this policy.
```

```

|scope
|link:#scope[scope]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "engine": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "string"
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "name": "fp_policy_1",
  "scope": {
    "exclude_export_policies": [
      "string"
    ],
    "exclude_extension": [
      "string"
    ],
    "exclude_shares": [
      "string"
    ],
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": [
      "string"
    ],
    "include_extension": [

```



```

    "string"
  ],
  "include_shares": [
    "sh1",
    "share_cifs"
  ],
  "include_volumes": [
    "vol1",
    "vol_svm1"
  ]
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",

```

```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#fpolicy_engine_reference]

```

```
[.api-collapsible-fifth-title]
fpolicy_engine_reference

FPolicy external engine

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the FPolicy external engine.

|===

[#fpolicy_event_reference]
[.api-collapsible-fifth-title]
fpolicy_event_reference

FPolicy events

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|===
```

```
[#scope]
[.api-collapsible-fifth-title]
scope
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|exclude_export_policies
|array[string]
a|
```

```
|exclude_extension
|array[string]
a|
```

```
|exclude_shares
|array[string]
a|
```

```
|exclude_volumes
|array[string]
a|
```

```
|include_export_policies
|array[string]
a|
```

```
|include_extension
|array[string]
a|
```

```
|include_shares
|array[string]
a|
```

```
|include_volumes
|array[string]
a|
```

```
|===
```

```
[#error_arguments]
```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID17a05e50180a034fc9072a936eaa22ad]]
= Update an FPolicy configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/fpolicy/{svm.uuid}/policies/{name}`#
```

Introduced In: 9.6

Updates a particular FPolicy policy configuration for a specified SVM. PATCH can be used to enable or disable the policy. When enabling a policy, you must specify the policy priority. The policy priority of the policy is not required when disabling the policy. If the policy is enabled, the FPolicy policy engine cannot be modified.

== Related ONTAP commands

```
* `fpolicy policy modify`
* `fpolicy policy scope modify`
* `fpolicy enable`
* `fpolicy disable`
```

== Learn more

```
*
xref:{relative_path}protocols_fpolicy_svm.uuid_policies_endpoint_overview.
html[DOC /protocols/fpolicy/{svm.uuid}/policies]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
```

```

|Type
|In
|Required
|Description

|name
|string
|path
|True
a|

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

```

```

|===

```

```

== Request Body

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|enabled
|boolean
a|Specifies if the policy is enabled on the SVM or not. If no value is
mentioned for this field but priority is set, then this policy will be
enabled.

```

```

|engine
|link:#fpolicy_engine_reference[fpolicy_engine_reference]
a|FPolicy external engine

```

```

|events
|array[link:#fpolicy_event_reference[fpolicy_event_reference]]
a|

```

```

|mandatory
|boolean

```

a|Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances.

|priority

|integer

a|Specifies the priority that is assigned to this policy.

|scope

|link:#scope[scope]

a|

|===

.Example request

[%collapsible%closed]

====

[source,json,subs+=macros]

```
{
  "engine": {
    "name": "string"
  },
  "events": [
    "event_nfs_close",
    "event_open"
  ],
  "scope": {
    "exclude_export_policies": [
      "string"
    ],
    "exclude_extension": [
      "string"
    ],
    "exclude_shares": [
      "string"
    ],
    "exclude_volumes": [
      "vol1",
      "vol_svm1",
      "*"
    ],
    "include_export_policies": [
```



```

    "string"
  ],
  "include_extension": [
    "string"
  ],
  "include_shares": [
    "sh1",
    "share_cifs"
  ],
  "include_volumes": [
    "vol1",
    "vol_svm1"
  ]
}
}
====

```

== Response

Status: 200, Ok

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
|
| 9765026
| The priority must be specified when enabling the FPolicy policy
|
| 9765025
| Cannot disable an FPolicy policy when the priority is specified
|
| 9764899
| Cannot modify an FPolicy engine when the policy is enabled
|
| 9764899
| Deletion of a cluster policy is not supported
|
| 9764908
| An FPolicy policy is already enabled

```

```
| 9764907
| An FPolicy policy is already disabled

| 9765029
| An FPolicy was modified but disable/enable failed as the policy is
already disabled/enabled
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#fpolicy_engine_reference]
[.api-collapsible-fifth-title]
fpolicy_engine_reference
```

FPolicy external engine

```
[cols=3*,options=header]
|===
|Name
|Type
```

```
|Description

|name
|string
a|The name of the FPolicy external engine.
```

```
|===
```

```
[#fpolicy_event_reference]
[.api-collapsible-fifth-title]
fpolicy_event_reference
```

```
FPolicy events
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|===
```

```
[#scope]
[.api-collapsible-fifth-title]
scope
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|exclude_export_policies
|array[string]
a|
```

```
|exclude_extension
|array[string]
a|
```

```
|exclude_shares  
|array[string]  
a|
```

```
|exclude_volumes  
|array[string]  
a|
```

```
|include_export_policies  
|array[string]  
a|
```

```
|include_extension  
|array[string]  
a|
```

```
|include_shares  
|array[string]  
a|
```

```
|include_volumes  
|array[string]  
a|
```

```
|===
```

```
[#fpolicy_policy]  
[.api-collapsible-fifth-title]  
fpolicy_policy
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|Specifies if the policy is enabled on the SVM or not. If no value is  
mentioned for this field but priority is set, then this policy will be  
enabled.
```

```
|engine
```

```
|link:#fpolicy_engine_reference[fpolicy_engine_reference]
```

a|FPolicy external engine

|events

|array[link:#fpolicy_event_reference[fpolicy_event_reference]]

a|

|mandatory

|boolean

a|Specifies what action to take on a file access event in a case when all primary and secondary servers are down or no response is received from the FPolicy servers within a given timeout period. When this parameter is set to true, file access events will be denied under these circumstances.

|priority

|integer

a|Specifies the priority that is assigned to this policy.

|scope

|link:#scope[scope]

a|

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= View NFS connected clients

:leveloffset: +1

```

```
[[IDba8aa81561995875fc8a8bcfbe616ada]]
= Protocols NFS connected-clients endpoint overview
```

`== Overview`

ONTAP connected clients show functionality is mainly used to provide a list of currently connected NFS clients. It also provides a potential list of other NFS clients that can be connected but are currently idle.

The following are details of the fields retrieved for the Connected Clients GET API:

node.name: The node name hosting this record; basically the node hosting the "server_ip".

node.uuid: The node UUID hosting this record; basically the node hosting the "server_ip".

svm.name: The svm name to which the "server_ip" belongs to.

svm.uuid: The svm uuid to which the "server_ip" belongs to.

server_ip: All clients that are connected to this interface are displayed in rows.

client_ip: The IP address of the client that is connected to the interface.

volume.name: The name of the volume the client is accessing.

volume.uuid: The UUID of the volume the client is accessing.

protocol: The NFS protocol version over which client is accessing the volume.

idle_duration: The time elapsed since the last request was sent by the client for this volume.

local_request_count: A counter that tracks requests that are sent to the volume with fast-path to local node.

remote_request_count: A counter that tracks requests that are sent to the volume with slow-path to remote node.

`== Example`

`=== Retrieves connected client information`

`----`

`# The API:`

`GET /protocols/nfs/connected-clients`

`# The call:`

```
curl -X GET "https://<cluster-mgmt-ip>/api/protocols/nfs/connected-clients?return_timeout=15&return_records=true" -H "accept: application/json"
```

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "c642db55-b8d0-11e9-9ad1-0050568e8480",
        "name": "vs1"
      },
      "node": {
        "uuid": "cc282893-b82f-11e9-a3ad-0050568e8480",
        "name": "vsim1"
      },
      "server_ip": "10.140.72.214",
      "client_ip": "10.140.137.57",
      "volume": {
        "name": "rvol1",
        "uuid": "c6bbc6f2-b8d0-11e9-9ad1-0050568e8480"
      },
      "protocol": "nfs4"
    },
    {
      "svm": {
        "uuid": "c642db55-b8d0-11e9-9ad1-0050568e8480",
        "name": "vs1"
      },
      "node": {
        "uuid": "cc282893-b82f-11e9-a3ad-0050568e8480",
        "name": "vsim1"
      },
      "server_ip": "10.140.72.214",
      "client_ip": "10.140.137.57",
      "volume": {
        "name": "vol1",
        "uuid": "d28d1999-b8d0-11e9-9ad1-0050568e8480"
      },
      "protocol": "nfs3"
    },
    {
      "svm": {
        "uuid": "c642db55-b8d0-11e9-9ad1-0050568e8480",
        "name": "vs1"
      },

```



```

"node": {
  "uuid": "cc282893-b82f-11e9-a3ad-0050568e8480",
  "name": "vsim1"
},
"server_ip": "10.140.72.214",
"client_ip": "10.140.137.57",
"volume": {
  "name": "vol1",
  "uuid": "d28d1999-b8d0-11e9-9ad1-0050568e8480"
},
"protocol": "nfs4"
}],
"num_records": 3
}
----

```

[[IDb0f73c7d7b2f8ecc2683a7128d00c0ea]]
= Retrieve the NFS configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nfs/connected-clients`#

Introduced In: 9.7

Retrieves the NFS configuration of SVMs.

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

local_request_count
integer
query
False
a Filter by local_request_count

```
|client_ip
|string
|query
|False
a|Filter by client_ip
```

```
|remote_request_count
|integer
|query
|False
a|Filter by remote_request_count
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|node.name
|string
|query
|False
a|Filter by node.name
```

```
|node.uuid
|string
|query
|False
a|Filter by node.uuid
```

```
|server_ip
|string
|query
```

```
|False
a|Filter by server_ip

|volume.uuid
|string
|query
|False
a|Filter by volume.uuid

|volume.name
|string
|query
|False
a|Filter by volume.name

|idle_duration
|string
|query
|False
a|Filter by idle_duration

|protocol
|string
|query
|False
a|Filter by protocol

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_timeout
```

```

|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|num_records
|integer
a|Number of records

|records
|array[link:#nfs_clients[nfs_clients]]
a|

|===

```

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "client_ip": "string",
      "idle_duration": "P4DT84H30M5S",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "protocol": "nfs3",
      "server_ip": "string",
      "svm": {

```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {

```

```

    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid

```



```

|string
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#volume]
[.api-collapsible-fifth-title]
volume

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|name
|string
a|The name of the volume.

|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6

|===

[#nfs_clients]
[.api-collapsible-fifth-title]
nfs_clients

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|client_ip
|string
a|Specifies IP address of the client.

|idle_duration
|string
a|Specifies an ISO-8601 format of date and time to retrieve the idle time
duration in hours, minutes, and seconds format.

|local_request_count
|integer
a|A counter that tracks requests that are sent to the volume with fast-
path to local node.

```

```

|node
|link:#node[node]
a|

|protocol
|string
a|The NFS protocol version over which client is accessing the volume. The
following values are supported:

* nfs - All NFS versions are considered
* nfs3 - NFS version 3 protocol
* nfs4 - NFS version 4 protocol
* nfs4.1 - NFS version 4 minor version 1 protocol

|remote_request_count
|integer
a|A counter that tracks requests that are sent to the volume with slow-
path to remote node.

|server_ip
|string
a|Specifies the IP address of the server.

|svm
|link:#svm[svm]
a|

|volume
|link:#volume[volume]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

```

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage NFS export policies

:leveloffset: +1

[[ID77479fe97530348758e19e750bb17860]]
= Protocols NFS export-policies endpoint overview

== Export Policies

=== 1) Retrieve the export policy details

'''

----

# The API:
GET /api/protocols/nfs/export-policies

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/export-policies"
----

'''

=== 2) Create an export policy for an SVM

'''

----

# The API:
POST /api/protocols/nfs/export-policies

# The call:
curl -d "@test_post_policy_single_rule.txt" -X POST "https://<mgmt-

```

```
ip>/api/protocols/nfs/export-policies"
```

```
test_post_policy_single_rule.txt(body):
```

```
{
  "name": "P1",
  "rules":[
    {
      "clients": [
        {
          "match": "host1"
        }
      ],
      "ro_rule": [
        "krb5"
      ],
      "rw_rule": [
        "ntlm"
      ],
      "anonymous_user": "anon1"
    },
    {
      "clients": [
        {
          "match": "host2"
        }
      ],
      "ro_rule": [
        "sys"
      ],
      "rw_rule": [
        "ntlm"
      ],
      "superuser": [
        "any"
      ]
    }
  ]
}
```

```
----
```

```
'''
```

```
=== 3) Update an export policy for an SVM
```

```
'''
```

```
----
```

```

# The API:
PATCH /api/protocols/nfs/export-policies/{policy.id}

# The call:
curl -d "@test_patch_policy.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/export-policies/8589934594"
test_patch_policy.txt (body):
{
  "name": "S1",
  "rules": [
    {
      "clients": [
        {
          "match": "host4"
        }
      ],
      "ro_rule": [
        "krb5"
      ],
      "rw_rule": [
        "ntlm"
      ]
    }
  ]
}
-----

'''

=== 4) Delete an export policy for an SVM

'''

-----

# The API:
DELETE /api/protocols/nfs/export-policies/{policy.id}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/nfs/export-
policies/8589934594"
-----

'''

```

== Export Rules

=== 1) Retrieve the export policy rule details for an export policy

'''

The API:

GET /api/protocols/nfs/export-policies/{policy.id}/rules

The call:

curl -X GET "https://<mgmt-ip>/api/protocols/nfs/export-policies/8589934595/rules"

'''

=== 2) Create an export policy rule for an export policy

'''

The API:

POST /api/protocols/nfs/export-policies/{policy.id}/rules

The call:

curl -d "@test_patch_export_rule.txt" -X POST "https://<mgmt-ip>/api/protocols/nfs/export-policies/8589934595/rules"

test_patch_export_rule.txt (body):

{

"clients": [

{

"match": "host2"

}

],

"ro_rule": [

"sys"

],

"rw_rule": [

"ntlm"

]

}

'''

=== 3) Update an export policy rule for an export policy

'''

The API:

PATCH /api/protocols/nfs/export-policies/{policy.id}/rules/{index}

The call:

curl -d "@test_patch_export_rule.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/export-policies/8589934595/rules/5?new_index=10"

test_patch_export_rule.txt(body):

```
{  
  "clients": [  
    {  
      "match": "host4"  
    }  
  ],  
  "ro_rule": [  
    "sys"  
  ],  
  "rw_rule": [  
    "krb5"  
  ]  
}
```

'''

=== 4) Delete an export policy rule for an export policy

'''

The API:

DELETE /api/protocols/nfs/export-policies/{policy.id}/rules/{index}

The call:

curl -X DELETE "https://<mgmt-ip>/api/protocols/nfs/export-
policies/8589934595/rules/15"

'''

== Export Clients

=== 1) Retrieve the export client matches of an export policy rule

'''

The API:

GET /api/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients

The call:

curl -X GET "https://<mgmt-ip>/api/protocols/nfs/export-policies/8589934593/rules/2/clients"

'''

=== 2) Add an export client match to an export policy rule

'''

The API:

POST /api/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients

The call:

curl -d "@add_client_match.txt" -X POST "https://<mgmt-ip>/api/protocols/nfs/export-policies/8589934593/rules/1/clients"
add_client_match.txt (body):

```
{  
  "match" : "host4"  
}
```

'''

=== 3) Delete an export client match from an export policy rule

'''

The API:

DELETE /api/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients/{match}

```

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/nfs/export-
policies/8589934593/rules/1/clients/host1,host2"
----

'''

[[ID9236618836cf1194c157791f69e9c7a1]]
= Retrieve export policies

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/export-policies`#

*Introduced In:* 9.6

Retrieves export policies.

== Related ONTAP commands

* `vserver export-policy show`
* `vserver export-policy rule show`

== Learn more

* xref:{relative_path}protocols_nfs_export-
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|id
|integer
|query

```

```
|False
a|Filter by id

|name
|string
|query
|False
a|Filter by name

|rules.anonymous_user
|string
|query
|False
a|Filter by rules.anonymous_user

|rules.index
|integer
|query
|False
a|Filter by rules.index

|rules.clients.match
|string
|query
|False
a|Filter by rules.clients.match

|rules.rw_rule
|string
|query
|False
a|Filter by rules.rw_rule

|rules.protocols
|string
|query
|False
a|Filter by rules.protocols

|rules.superuser
```

```
|string
|query
|False
a|Filter by rules.superuser
```

```
|rules.ro_rule
|string
|query
|False
a|Filter by rules.ro_rule
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

|return_timeout
|integer
|query
|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|The number of export policy records

```
|records
|array[link:#export_policy[export_policy]]
a|

|===
```

.Example response

[%collapsible%closed]

=====

```
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": 0,
      "name": "string",
      "rules": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "anonymous_user": "string",
          "clients": [
            {
              "match": "0.0.0.0/0"
            }
          ],
          "index": 0,
          "protocols": [
            "string"
          ],
        }
      ]
    }
  ]
}
```

```

        "ro_rule": [
            "string"
        ],
        "rw_rule": [
            "string"
        ],
        "superuser": [
            "string"
        ]
    }
],
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====

```



```

[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#export_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|===
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

```

- * As a hostname; for instance, host1
- * As an IPv4 address; for instance, 10.1.12.24
- * As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
- * As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24
- * As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8ble:b255:4071::/64
- * As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- * As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng
- * As a domain name preceded by the . character; for instance, .example.com

|===

```
[#export_rule]
[.api-collapsible-fifth-title]
export_rule
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|anonymous_user
|string
a|User ID To Which Anonymous Users Are Mapped.
```

```
|clients
|array[link:#export_client[export_client]]
a|Array of client matches
```

```
|index
|integer
a|Index of the rule within the export policy.
```

```

|protocols
|array[string]
a|

|ro_rule
|array[string]
a|Authentication flavors that the read-only access rule governs

|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs

|superuser
|array[string]
a|Authentication flavors that the superuser security type governs

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

```

```

[#export_policy]
[.api-collapsible-fifth-title]
export_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|Export Policy ID

|name
|string
a|Export Policy Name

|rules
|array[link:#export_rule[export_rule]]
a|Rules of the Export Policy.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[IDf9d3920a063b30ac2599874728f7cfc7]]
```

```
= Create an export policy
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-  
block]#`/protocols/nfs/export-policies`#
```

```
*Introduced In:* 9.6
```

Creates an export policy. An SVM can have any number of export policies to define rules for which clients can access data exported by the SVM. A policy with no rules prohibits access.

```
== Required properties
```

```
* `svm.uuid` or `svm.name` - Existing SVM in which to create an export  
policy.
```

```
* `name` - Name of the export policy.
```

```
== Recommended optional properties
```

```
* `rules` - Rule(s) of an export policy. Used to create the export rule  
and populate the export policy with export rules in a single request.
```

```
== Related ONTAP commands
```

```
* `vserver export-policy create`
```

```
* `vserver export-policy rule create`
```

```
== Learn more
```

```
* xref:{relative_path}protocols_nfs_export-  
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|==
```

```
|Name
```

```
|Type
```

```
|In
```

```

|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Export Policy ID

|name
|string
a|Export Policy Name

|rules
|array[link:#export_rule[export_rule]]
a|Rules of the Export Policy.

|svm
|link:#svm[svm]
a|

|===

.Example request
[%collapsible%closed]

```



```

=====
[source,json,subs=+macros]
{
  "id": 0,
  "name": "string",
  "rules": [
    {
      "anonymous_user": "string",
      "clients": [
        {
          "match": "0.0.0.0/0"
        }
      ],
      "index": 0,
      "protocols": [
        "string"
      ],
      "ro_rule": [
        "string"
      ],
      "rw_rule": [
        "string"
      ],
      "superuser": [
        "string"
      ]
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

== Response

```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|num_records
|integer
a|The number of export policy records

|records
|array[link:#export_policy[export_policy]]
a|

|===

```

```

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]

```

```

{
  "records": [
    {
      "id": 0,
      "name": "string",
      "rules": [
        {
          "anonymous_user": "string",
          "clients": [
            {
              "match": "0.0.0.0/0"
            }
          ],
          "index": 0,
          "protocols": [
            "string"
          ],
          "ro_rule": [
            "string"
          ],
          "rw_rule": [
            "string"
          ],
          "superuser": [
            "string"
          ]
        }
      ],
      "svm": {
        "name": "svm1",

```

```

        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
}
]
}
====

```

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

1703952	
---------	--

Invalid ruleset name provided. No spaces allowed in a ruleset name

1703954	
---------	--

Export policy does not exist

1704049	
---------	--

Invalid clientmatch: clientmatch lists require an effective cluster version of Data ONTAP 9.0 or later. Upgrade all nodes to Data ONTAP 9.0 or above to use features that operate on lists of clientmatch strings in export-policy rules

1704055	
---------	--

Export policies are only supported for data Vservers

3277000	
---------	--

Upgrade all nodes to Data ONTAP 9.0.0 or above to use krb5p as a security flavor in export-policy rules

3277083	
---------	--

User ID is not valid. Enter a value for User ID from 0 to 4294967295

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|====
|Name
|Type
|Description

|href
|string
a|

|====

[#\_links]
[.api-collapsible-fifth-title]
_links
[#export\_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|====
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
    following formats:

* As a hostname; for instance, host1
* As an IPv4 address; for instance, 10.1.12.24
* As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1
* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8b1e:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0

```

```

* As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
* As a domain name preceded by the . character; for instance, .example.com

|===

[#export_rule]
[.api-collapsible-fifth-title]
export_rule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|anonymous_user
|string
a|User ID To Which Anonymous Users Are Mapped.

|clients
|array[link:#export_client[export_client]]
a|Array of client matches

|index
|integer
a|Index of the rule within the export policy.

|protocols
|array[string]
a|

|ro_rule
|array[string]
a|Authentication flavors that the read-only access rule governs

|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs

```

```

|superuser
|array[string]
a|Authentication flavors that the superuser security type governs

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#export_policy]
[.api-collapsible-fifth-title]
export_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Export Policy ID

|name

```

```
|string
a|Export Policy Name
```

```
|rules
|array[link:#export_rule[export_rule]]
a|Rules of the Export Policy.
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
[#export_client]
[.api-collapsible-fifth-title]
export_client
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|match
|string
```

```
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
    following formats:
```

- * As a hostname; for instance, host1
- * As an IPv4 address; for instance, 10.1.12.24
- * As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1
- * As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24
- * As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8b1e:b255:4071::/64
- * As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- * As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng
- * As a domain name preceded by the . character; for instance, .example.com

```
|===
```

```
[#export_client]
[.api-collapsible-fifth-title]
export_client
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|match
```

```
|string
```

```
a|Client Match Hostname, IP Address, Netgroup, or Domain.
```

```
You can specify the match as a string value in any of the
    following formats:
```

```
* As a hostname; for instance, host1
```

```
* As an IPv4 address; for instance, 10.1.12.24
```

```
* As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
```

```
* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
```

```
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
```

```
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
```

```
* As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
```

```
* As a domain name preceded by the . character; for instance, .example.com
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

```
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```



```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```

[[ID80e096794b3e1a27694bebc6cd02b6a5]]
= Delete an export policy

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/nfs/export-policies/{id}`#

*Introduced In:* 9.6

Deletes an export policy.

== Related ONTAP commands

* `vserver export-policy delete`

== Learn more

* xref:{relative_path}protocols_nfs_export-
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|id
|integer
|path
|True
a|Export Policy ID

|===

== Response

```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

```
|===  
| Error Code | Description  
  
| 1703944  
| Failed to delete rule  
  
| 1703945  
| Ruleset is in use by a volume. It cannot be deleted until all volumes  
that refer to it are first deleted  
  
| 1703946  
| Cannot determine if the ruleset is in use by a volume. It cannot be  
deleted until all volumes that refer to it are first deleted  
  
| 1703947  
| Cannot delete default ruleset. This ruleset will be deleted when the  
owning Vserver is deleted  
  
| 1703952  
| Invalid ruleset name provided. No spaces are allowed in a ruleset name  
  
| 1703953  
| This ruleset is in use by a qtree export policy. It cannot be deleted  
until all qtree policies that refer to it are first deleted  
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID08e65157dfb7a59d15388213a9af3295]]
= Retrieve an export policy

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/export-policies/{id}`#

```

Introduced In: 9.6

Retrieves an export policy.

== Related ONTAP commands

```
* `vserver export-policy show`  
* `vserver export-policy rule show`
```

== Learn more

* xref:{relative_path}protocols_nfs_export-policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|id  
|integer  
|path  
|True  
a|Export Policy ID
```

```
|fields  
|array[string]  
|query  
|False  
a|Specify the fields to return.
```

```
|===
```

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
a|Export Policy ID

|name
|string
a|Export Policy Name

|rules
|array[link:#export_rule[export_rule]]
a|Rules of the Export Policy.

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": 0,
  "name": "string",
  "rules": [
    {
      "_links": {

```

```

        "self": {
            "href": "/api/resourcelink"
        },
    },
    "anonymous_user": "string",
    "clients": [
        {
            "match": "0.0.0.0/0"
        }
    ],
    "index": 0,
    "protocols": [
        "string"
    ],
    "ro_rule": [
        "string"
    ],
    "rw_rule": [
        "string"
    ],
    "superuser": [
        "string"
    ]
}
],
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type

```



```

|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href

```

```

|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#export_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|===
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

* As a hostname; for instance, host1
* As an IPv4 address; for instance, 10.1.12.24
* As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,

```

10.1.16.0/255.255.255.0

* As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng

* As a domain name preceded by the . character; for instance, .example.com

|===

[#export_rule]

[.api-collapsible-fifth-title]

export_rule

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|anonymous_user

|string

a|User ID To Which Anonymous Users Are Mapped.

|clients

|array[link:#export_client[export_client]]

a|Array of client matches

|index

|integer

a|Index of the rule within the export policy.

|protocols

|array[string]

a|

|ro_rule

|array[string]

a|Authentication flavors that the read-only access rule governs

```

|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs

|superuser
|array[string]
a|Authentication flavors that the superuser security type governs

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

```

```
//end collapsible .Definitions block
====
```

```
[[ID874c0268a21772393c0fa39bb4fb2a8f]]
```

```
= Update export policy properties
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-  
block]#`/protocols/nfs/export-policies/{id}`#
```

```
*Introduced In:* 9.6
```

Updates the properties of an export policy to change an export policy name or replace all export policy rules.

```
== Related ONTAP commands
```

```
* `vserver export-policy rename`  
* `vserver export-policy rule delete`  
* `vserver export-policy rule create`
```

```
== Learn more
```

```
* xref:{relative_path}protocols_nfs_export-  
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]
```

```
== Parameters
```

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|id  
|integer  
|path  
|True  
a|Export Policy ID
```

```

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|integer
a|Export Policy ID

|name
|string
a|Export Policy Name

|rules
|array[link:#export_rule[export_rule]]
a|Rules of the Export Policy.

|svm
|link:#svm[svm]
a|

|===

```

```

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "id": 0,
  "name": "string",
  "rules": [
    {
      "anonymous_user": "string",
      "clients": [
        {
          "match": "0.0.0.0/0"

```

```

    }
    ],
    "index": 0,
    "protocols": [
        "string"
    ],
    "ro_rule": [
        "string"
    ],
    "rw_rule": [
        "string"
    ],
    "superuser": [
        "string"
    ]
}
],
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

```

== Response

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 1703950
| Failed to rename ruleset
| 1703952
| Invalid ruleset name provided. No spaces are allowed in a ruleset name
|===

```



```
== Definitions
```

```
[.api-def-first-level]  
.See Definitions  
[%collapsible%closed]  
//Start collapsible Definitions block
```

```
====
```

```
[#href]  
[.api-collapsible-fifth-title]  
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|href  
|string  
a|
```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links  
[#export_client]  
[.api-collapsible-fifth-title]  
export_client
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|match  
|string
```

a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

- * As a hostname; for instance, host1
- * As an IPv4 address; for instance, 10.1.12.24
- * As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1

- * As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24
- * As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8ble:b255:4071::/64
- * As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- * As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng
- * As a domain name preceded by the . character; for instance, .example.com

|===

```
[#export_rule]
[.api-collapsible-fifth-title]
export_rule
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|anonymous_user
```

```
|string
```

```
a|User ID To Which Anonymous Users Are Mapped.
```

```
|clients
```

```
|array[link:#export_client[export_client]]
```

```
a|Array of client matches
```

```
|index
```

```
|integer
```

```
a|Index of the rule within the export policy.
```

```
|protocols
```

```
|array[string]
```

```
a|
```

```
|ro_rule
```

```
|array[string]
```

```
a|Authentication flavors that the read-only access rule governs
```

```
|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs
```

```
|superuser
|array[string]
a|Authentication flavors that the superuser security type governs
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#export_policy]
[.api-collapsible-fifth-title]
export_policy
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```

|id
|integer
a|Export Policy ID

|name
|string
a|Export Policy Name

|rules
|array[link:#export_rule[export_rule]]
a|Rules of the Export Policy.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]

```

```

error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====


[[IDccefd787d3faaad940a111c0165d1471]]
= Retrieve export policy rules


[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/export-policies/{policy.id}/rules`#


*Introduced In:* 9.6


Retrieves export policy rules.


== Related ONTAP commands

```

```
* `vserver export-policy rule show`
```

== Learn more

```
* xref:{relative_path}protocols_nfs_export-  
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]
```

== Parameters

```
[cols=5*,options=header]  
|==
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|policy.id  
|integer  
|path  
|True  
a|Export Policy ID
```

```
|anonymous_user  
|string  
|query  
|False  
a|Filter by anonymous_user
```

```
|index  
|integer  
|query  
|False  
a|Filter by index
```

```
|clients.match  
|string  
|query  
|False  
a|Filter by clients.match
```

```
|rw_rule
|string
|query
|False
a|Filter by rw_rule
```

```
|protocols
|string
|query
|False
a|Filter by protocols
```

```
|superuser
|string
|query
|False
a|Filter by superuser
```

```
|ro_rule
|string
|query
|False
a|Filter by ro_rule
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
```

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of Export Rule records


```
|records
|array[link:#export_rule[export_rule]]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anonymous_user": "string",
      "clients": [
        {
          "match": "0.0.0.0/0"
        }
      ],
      "index": 0,
      "protocols": [
        "string"
      ],
      "ro_rule": [
        "string"
      ],
      "rw_rule": [
        "string"
      ],
      "superuser": [
```

```

        "string"
    ]
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]

```

```

.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#export_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|===
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

* As a hostname; for instance, host1
* As an IPv4 address; for instance, 10.1.12.24
* As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
* As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
* As a domain name preceded by the . character; for instance, .example.com

|===

[#export_rule]
[.api-collapsible-fifth-title]

```

```

export_rule

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|anonymous_user
|string
a|User ID To Which Anonymous Users Are Mapped.

|clients
|array[link:#export_client[export_client]]
a|Array of client matches

|index
|integer
a|Index of the rule within the export policy.

|protocols
|array[string]
a|

|ro_rule
|array[string]
a|Authentication flavors that the read-only access rule governs

|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs

|superuser
|array[string]
a|Authentication flavors that the superuser security type governs

|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[IDdffc09ef5b8643b4b5d202df83a77d0a]]

= Create an export policy rule

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/protocols/nfs/export-policies/{policy.id}/rules`#

Introduced In: 9.6

Creates an export policy rule.

== Required properties

* `policy.id` - Existing export policy for which to create an export rule.

* `clients.match` - List of clients (hostnames, ipaddresses, netgroups, domains) to which the export rule applies.

* `ro_rule` - Used to specify the security type for read-only access to volumes that use the export rule.

* `rw_rule` - Used to specify the security type for read-write access to volumes that use the export rule.

== Default property values

If not specified in POST, the following default property values are assigned:

* `protocols` - `_any_`

* `anonymous_user` - `_none_`

* `superuser` - `_any_`

== Related ONTAP commands

* `vserver export-policy rule create`

== Learn more

* xref:{relative_path}protocols_nfs_export-policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|policy.id

|integer

|path

|True

a|Export Policy ID

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description


```

|anonymous_user
|string
a|User ID To Which Anonymous Users Are Mapped.

|clients
|array[link:#export_client[export_client]]
a|Array of client matches

|index
|integer
a|Index of the rule within the export policy.

|protocols
|array[string]
a|

|ro_rule
|array[string]
a|Authentication flavors that the read-only access rule governs

|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs

|superuser
|array[string]
a|Authentication flavors that the superuser security type governs

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "anonymous_user": "string",
  "clients": [
    {
      "match": "0.0.0.0/0"
    }
  ]
}

```

```

],
"index": 0,
"protocols": [
  "string"
],
"ro_rule": [
  "string"
],
"rw_rule": [
  "string"
],
"superuser": [
  "string"
]
}
====

== Response

```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of Export Rule records

|records
|array[link:#export_rule[export_rule]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {

```

```

    "anonymous_user": "string",
    "clients": [
      {
        "match": "0.0.0.0/0"
      }
    ],
    "index": 0,
    "protocols": [
      "string"
    ],
    "ro_rule": [
      "string"
    ],
    "rw_rule": [
      "string"
    ],
    "superuser": [
      "string"
    ]
  }
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 1703954
| Export policy does not exist

| 1704036
| Invalid clientmatch:  missing domain name

| 1704037
| Invalid clientmatch:  missing network name

| 1704038
| Invalid clientmatch:  missing netgroup name

| 1704039

```

```

| Invalid clientmatch

| 1704040
| Invalid clientmatch: address bytes masked out by netmask are non-zero

| 1704041
| Invalid clientmatch: address bytes masked to zero by netmask

| 1704042
| Invalid clientmatch: too many bits in netmask

| 1704043
| Invalid clientmatch: invalid netmask

| 1704044
| Invalid clientmatch: invalid characters in host name

| 1704045
| Invalid clientmatch: invalid characters in domain name

| 1704050
| Invalid clientmatch: clientmatch list contains a duplicate string.
Duplicate strings in a clientmatch list are not supported

| 1704051
| Warning: Not adding any new strings to the clientmatch field for
ruleindex. All of the match strings are already in the clientmatch list

| 1704064
| Clientmatch host name too long

| 1704065
| Clientmatch domain name too long

| 3277000
| Upgrade all nodes to Data ONTAP 9.0.0 or above to use krb5p as a
security flavor in export-policy rules

| 3277083
| User ID is not valid. Enter a value for User ID from 0 to 4294967295
|===

== Definitions

[.api-def-first-level]
.See Definitions

```

```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#export_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|===
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

* As a hostname; for instance, host1
* As an IPv4 address; for instance, 10.1.12.24
* As an IPv6 address; for instance, fd20:8b1e:b255:4071::100:1
* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8b1e:b255:4071::/64
```

- * As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- * As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng
- * As a domain name preceded by the . character; for instance, .example.com

|===

[#export_rule]
[.api-collapsible-fifth-title]
export_rule

[cols=3*,options=header]

|===

|Name
|Type
|Description

|anonymous_user
|string
a|User ID To Which Anonymous Users Are Mapped.

|clients
|array[link:#export_client[export_client]]
a|Array of client matches

|index
|integer
a|Index of the rule within the export policy.

|protocols
|array[string]
a|

|ro_rule
|array[string]
a|Authentication flavors that the read-only access rule governs

|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs

```
|superuser
|array[string]
a|Authentication flavors that the superuser security type governs
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID375b32c24abebd38f008bab0790d3be2]]
= Delete an export policy rule

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/nfs/export-policies/{policy.id}/rules/{index}`#

*Introduced In:* 9.6

Deletes an export policy rule.

== Related ONTAP commands

* `vserver export-policy rule delete`

== Learn more

* xref:{relative_path}protocols_nfs_export-
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]

== Parameters

[cols=5*,options=header]

```



```

|===
|Name
|Type
|In
|Required
|Description

|policy.id
|integer
|path
|True
a|Export Policy ID

|index
|integer
|path
|True
a|Export Rule Index

|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 1703945
| Ruleset is in use by a volume. It cannot be deleted until all volumes
that refer to it are first deleted

| 1703946
| Cannot determine if the ruleset is in use by a volume. It cannot be
deleted until all volumes that refer to it are first deleted

| 1703954

```

```
| Export policy does not exist
```

```
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "error": {
```

```
    "arguments": [
```

```
      {
```

```
        "code": "string",
```

```
        "message": "string"
```

```
      }
```

```
    ],
```

```
    "code": "4",
```

```
    "message": "entry doesn't exist",
```

```
    "target": "uuid"
```

```
  }
```

```
}
```

```
=====
```

```
== Definitions
```

```
[.api-def-first-level]
```

```
.See Definitions
```

```
[%collapsible%closed]
```

```
//Start collapsible Definitions block
```

```
=====
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID0cc5313b5ee35f212777607e5b10e515]]
= Retrieve an export policy rule

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/export-policies/{policy.id}/rules/{index}`#

*Introduced In:* 9.6

Retrieves an export policy rule

== Related ONTAP commands

* `vserver export-policy rule show`

== Learn more

* xref:{relative_path}protocols_nfs_export-
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|policy.id
|integer
|path
|True

```

```

a|Export Policy ID

|index
|integer
|path
|True
a|Export Rule Index

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|anonymous_user
|string
a|User ID To Which Anonymous Users Are Mapped.

|clients
|array[link:#export_client[export_client]]
a|Array of client matches

|index
|integer
a|Index of the rule within the export policy.

```

```

|protocols
|array[string]
a|

|ro_rule
|array[string]
a|Authentication flavors that the read-only access rule governs

|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs

|superuser
|array[string]
a|Authentication flavors that the superuser security type governs

```

```
|===
```

```

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anonymous_user": "string",
  "clients": [
    {
      "match": "0.0.0.0/0"
    }
  ],
  "index": 0,
  "protocols": [
    "string"
  ],
  "ro_rule": [
    "string"
  ],

```

```

"rw_rule": [
  "string"
],
"superuser": [
  "string"
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#export_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|===
|Name

```



```

|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

* As a hostname; for instance, host1
* As an IPv4 address; for instance, 10.1.12.24
* As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
* As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
* As a domain name preceded by the . character; for instance, .example.com

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDef594b0b44ea55b77b3777e6b3a0f045]]
= Update the properties of an export policy rule

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/nfs/export-policies/{policy.id}/rules/{index}`#

*Introduced In:* 9.6

```

Updates the properties of an export policy rule to change an export policy rule's index or fields.

== Related ONTAP commands

- * ``vserver export-policy rule modify``
- * ``vserver export-policy rule setindex``

== Learn more

- * `xref:{relative_path}protocols_nfs_export-policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]`

== Parameters

[cols=5*,options=header]

|===

Name
Type
In
Required
Description

policy.id
integer
path
True
a Export Policy ID

index
integer
path
True
a Export Rule Index

new_index
integer
query
False
a New Export Rule Index

|===

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|anonymous_user
```

```
|string
```

```
a|User ID To Which Anonymous Users Are Mapped.
```

```
|clients
```

```
|array[link:#export_client[export_client]]
```

```
a|Array of client matches
```

```
|index
```

```
|integer
```

```
a|Index of the rule within the export policy.
```

```
|protocols
```

```
|array[string]
```

```
a|
```

```
|ro_rule
```

```
|array[string]
```

```
a|Authentication flavors that the read-only access rule governs
```

```
|rw_rule
```

```
|array[string]
```

```
a|Authentication flavors that the read/write access rule governs
```

```
|superuser
```

```
|array[string]
```

```
a|Authentication flavors that the superuser security type governs
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "anonymous_user": "string",
  "clients": [
    {
      "match": "0.0.0.0/0"
    }
  ],
  "index": 0,
  "protocols": [
    "string"
  ],
  "ro_rule": [
    "string"
  ],
  "rw_rule": [
    "string"
  ],
  "superuser": [
    "string"
  ]
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description
| 1703954
| Export policy does not exist
| 1704036
```

```
| Invalid clientmatch: missing domain name

| 1704037
| Invalid clientmatch: missing network name

| 1704038
| Invalid clientmatch: missing netgroup name

| 1704039
| Invalid clientmatch

| 1704040
| Invalid clientmatch: address bytes masked out by netmask are non-zero

| 1704041
| Invalid clientmatch: address bytes masked to zero by netmask

| 1704042
| Invalid clientmatch: too many bits in netmask

| 1704043
| Invalid clientmatch: invalid netmask

| 1704044
| Invalid clientmatch: invalid characters in host name

| 1704045
| Invalid clientmatch: invalid characters in domain name

| 1704050
| Invalid clientmatch: clientmatch list contains a duplicate string.
Duplicate strings in a clientmatch list are not supported

| 1704051
| Warning: Not adding any new strings to the clientmatch field for
ruleindex. All of the match strings are already in the clientmatch list

| 1704064
| Clientmatch host name too long

| 1704065
| Clientmatch domain name too long

| 3277000
| Upgrade all nodes to Data ONTAP 9.0.0 or above to use krb5p as a
security flavor in export-policy rules
```

```

| 3277083
| User ID is not valid. Enter a value for User ID from 0 to 4294967295
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#export_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|===
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

```

- * As a hostname; for instance, host1
- * As an IPv4 address; for instance, 10.1.12.24
- * As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
- * As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24
- * As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8ble:b255:4071::/64
- * As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- * As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng
- * As a domain name preceded by the . character; for instance, .example.com

|===

```
[#export_rule]
[.api-collapsible-fifth-title]
export_rule
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|anonymous_user
|string
a|User ID To Which Anonymous Users Are Mapped.
```

```
|clients
|array[link:#export_client[export_client]]
a|Array of client matches
```

```
|index
|integer
a|Index of the rule within the export policy.
```

```
|protocols
|array[string]
a|
```



```

|ro_rule
|array[string]
a|Authentication flavors that the read-only access rule governs


|rw_rule
|array[string]
a|Authentication flavors that the read/write access rule governs


|superuser
|array[string]
a|Authentication flavors that the superuser security type governs


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments


[cols=3*,options=header]
|===
|Name
|Type
|Description


|code
|string
a|Argument code


|message
|string
a|Message argument


|===


[#error]
[.api-collapsible-fifth-title]
error


[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

[[IDd5b2c535e3395dddddcb18e7ebd50b566]]
= Retrieve export policy rule clients

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients`#

*Introduced In:* 9.6

Retrieves export policy rule clients.

== Learn more

* xref:{relative_path}protocols_nfs_export-
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]

```

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|policy.id
|integer
|path
|True
a|Export Policy ID

|index
|integer
|path
|True
a|Export Rule Index

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of export rule client records


|records
|array[link:#export_client[export_client]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "match": "0.0.0.0/0"
    }
  ]
}
====

== Error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#export_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|===
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the

```

following formats:

- * As a hostname; for instance, host1
- * As an IPv4 address; for instance, 10.1.12.24
- * As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
- * As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24
- * As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8ble:b255:4071::/64
- * As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- * As a netgroup, with the netgroup name preceded by the @ character; for instance, @eng
- * As a domain name preceded by the . character; for instance, .example.com

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

|===

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====

[[IDada1662f5e6f1fb9f64258512daf24cc]]
= Create an export policy rule client


[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/nfs/export-policies/{policy.id}/rules/{index}/clients`#


*Introduced In:* 9.6


Creates an export policy rule client


== Required properties


* `policy.id` - Existing export policy that contains export policy rules
for the client being added.

```

* ``index`` - Existing export policy rule for which to create an export client.

* ``match`` - Base name for the export policy client.

== Related ONTAP commands

* ``vserver export-policy rule add-clientmatches``

== Learn more

* `xref:{relative_path}protocols_nfs_export-policies_endpoint_overview.html`[DOC /protocols/nfs/export-policies]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|policy.id

|integer

|path

|True

a|Export Policy ID

|index

|integer

|path

|True

a|Export Rule Index

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

```

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

* As a hostname; for instance, host1
* As an IPv4 address; for instance, 10.1.12.24
* As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
* As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
* As a domain name preceded by the . character; for instance, .example.com

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "match": "0.0.0.0/0"
}
=====

== Response

```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of export rule client records

|records
|array[link:#export_client[export_client]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "match": "0.0.0.0/0"
    }
  ]
}
====

== Error
```

Status: Default

```
ONTAP Error Response Codes

|===
| Error Code | Description

| 1703954
| Export policy does not exist
```

```
| 1704036
| Invalid clientmatch: missing domain name

| 1704037
| Invalid clientmatch: missing network name

| 1704038
| Invalid clientmatch: missing netgroup name

| 1704039
| Invalid clientmatch

| 1704040
| Invalid clientmatch: address bytes masked out by netmask are non-zero

| 1704041
| Invalid clientmatch: address bytes masked to zero by netmask

| 1704042
| Invalid clientmatch: too many bits in netmask

| 1704043
| Invalid clientmatch: invalid netmask

| 1704044
| Invalid clientmatch: invalid characters in host name

| 1704045
| Invalid clientmatch: invalid characters in domain name

| 1704050
| Invalid clientmatch: the clientmatch list contains a duplicate string.
Duplicate strings in a clientmatch list are not supported

| 1704051
| Warning: Not adding any new strings to the clientmatch field for
ruleindex. All of the match strings are already in the clientmatch list

| 1704064
| Clientmatch host name too long

| 1704065
| Clientmatch domain name too long
|===
```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#export_client]
[.api-collapsible-fifth-title]
export_client

[cols=3*,options=header]
|===
|Name
|Type
|Description

|match
|string
a|Client Match Hostname, IP Address, Netgroup, or Domain.
You can specify the match as a string value in any of the
following formats:

* As a hostname; for instance, host1
* As an IPv4 address; for instance, 10.1.12.24
* As an IPv6 address; for instance, fd20:8ble:b255:4071::100:1
* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
* As a netgroup, with the netgroup name preceded by the @ character; for
instance, @eng
* As a domain name preceded by the . character; for instance, .example.com

|===

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments

```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[ID55c1154ece3e3dale6fale547ed62a81]]
```

```
= Delete an export policy client
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/protocols/nfs/export-  
policies/{policy.id}/rules/{index}/clients/{match}`#
```

```
*Introduced In:* 9.6
```

```
Deletes an export policy client
```

```
== Related ONTAP commands
```

```
* `vserver export-policy rule remove-clientmatches`
```

```
== Learn more
```

```
* xref:{relative_path}protocols_nfs_export-  
policies_endpoint_overview.html[DOC /protocols/nfs/export-policies]
```



```
== Parameters
```

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|policy.id  
|integer  
|path  
|True  
a|Export Policy ID
```

```
|index  
|integer  
|path  
|True  
a|Export Rule Index
```

```
|match  
|string  
|path  
|True  
a|Export Client Match
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes
```

```
|===
```

Error Code	Description
1703954	Export policy does not exist
1704036	Invalid clientmatch: missing domain name
1704037	Invalid clientmatch: missing network name
1704038	Invalid clientmatch: missing netgroup name
1704039	Invalid clientmatch
1704040	Invalid clientmatch: address bytes masked out by netmask are non-zero
1704041	Invalid clientmatch: address bytes masked to zero by netmask
1704042	Invalid clientmatch: too many bits in netmask
1704043	Invalid clientmatch: invalid netmask
1704044	Invalid clientmatch: invalid characters in host name
1704045	Invalid clientmatch: invalid characters in domain name
1704050	Invalid clientmatch: the clientmatch list contains a duplicate string. Duplicate strings in a clientmatch list are not supported
1704052	Warning: Not removing any strings from the clientmatch field for ruleindex. None of the match strings were found in the clientmatch list
1704064	Clientmatch host name too long
1704065	

```

| Clientmatch domain name too long
|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#error_arguments]
[.api-collapsible-fifth-title]

```

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

:leveloffset: -1

= View and update Kerberos interfaces

:leveloffset: +1

[[IDc9eb2cb566532c5805cb2308b9335f98]]
= Protocols NFS Kerberos interfaces endpoint overview

== Examples

=== Retrieving the Kerberos interface configuration details

----

# The API:
GET /api/protocols/nfs/kerberos/interfaces

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/kerberos/interfaces"
----

=== Updating the Kerberos interface configuration

----

# The API:
PATCH /api/protocols/nfs/kerberos/interfaces/{interface.uuid}

# The call:
curl -d "@test_patch_kerb_interface.txt" -X PATCH "https://<mgmt-

```

```
ip>/api/protocols/nfs/kerberos/interfaces/e62936de-7342-11e8-9eb4-0050568be2b7"
test_patch_kerb_interface.txt (body):
{
  "enabled" : "true",
  "spn": "nfs/datalif1-vsim3-d1.sim.netapp.com@NFS-NSR-W01.RTP.NETAPP.COM",
  "user" : "administrator",
  "password" : "Hello123!"
}
----
```

```
[[ID367651f35c5906845828359cdbf8aaa0]]
= Retrieve Kerberos interfaces
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nfs/kerberos/interfaces`#
```

Introduced In: 9.6

Retrieves Kerberos interfaces.

== Related ONTAP commands

* `vserver nfs kerberos interface show`

== Learn more

*

xref:{relative_path}protocols_nfs_kerberos_interfaces_endpoint_overview.html[DOC /protocols/nfs/kerberos/interfaces]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|svm.uuid  
|string  
|query  
|False  
a|Filter by svm.uuid
```

```
|svm.name  
|string  
|query  
|False  
a|Filter by svm.name
```

```
|encryption_types  
|string  
|query  
|False  
a|Filter by encryption_types
```

```
|spn  
|string  
|query  
|False  
a|Filter by spn
```

```
|enabled  
|boolean  
|query  
|False  
a|Filter by enabled
```

```
|interface.name  
|string  
|query  
|False  
a|Filter by interface.name
```

```
|interface.uuid  
|string  
|query  
|False  
a|Filter by interface.uuid
```

```
|interface.ip.address
|string
|query
|False
a|Filter by interface.ip.address
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1
* Max value: 120
* Min value: 0


```

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#kerberos_interface[kerberos_interface]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  }
}

```

```

    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "encryption_types": [
        "string"
      ],
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "spn": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#ip]
[.api-collapsible-fifth-title]

```

ip

IP information

[cols=3*,options=header]

|===

|Name

|Type

|Description

|address

|string

a|IPv4 or IPv6 address

|===

[#interface]

[.api-collapsible-fifth-title]

interface

Network interface

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|ip

|link:#ip[ip]

a|IP information

|name

|string

a|The name of the interface.

|uuid

```
|string
a|The UUID that uniquely identifies the interface.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#kerberos_interface]
[.api-collapsible-fifth-title]
kerberos_interface
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|enabled
|boolean
a|Specifies if Kerberos is enabled.
```

```
|encryption_types
|array[string]
a|
```

```
|interface
|link:#interface[interface]
a|Network interface
```

```
|spn
|string
a|Service principal name. Valid in PATCH.
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDa71956b039c112b0c580318227775804]]
= Retrieve a Kerberos interface

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/kerberos/interfaces/{interface.uuid}`#

```


Introduced In: 9.6

Retrieves a Kerberos interface.

== Related ONTAP commands

* `vserver nfs kerberos interface show`

== Learn more

*

xref:{relative_path}protocols_nfs_kerberos_interfaces_endpoint_overview.html[DOC /protocols/nfs/kerberos/interfaces]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|interface.uuid

|string

|path

|True

a|Network interface UUID

* Introduced in: 9.7

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|Specifies if Kerberos is enabled.

|encryption_types
|array[string]
a|

|interface
|link:#interface[interface]
a|Network interface

|spn
|string
a|Service principal name. Valid in PATCH.

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },

```

```

"encryption_types": [
  "string"
],
"interface": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7"
  },
  "name": "lif1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"spn": "string",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#ip]
[.api-collapsible-fifth-title]
ip
```

IP information

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|===
```

```
[#interface]
[.api-collapsible-fifth-title]
interface
```

Network interface

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```

|_links
|link:#_links[_links]
a|

|ip
|link:#ip[ip]
a|IP information

|name
|string
a|The name of the interface.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

=====

[[ID620907128ffdf2f4a36c9a438e8b7e41]]

= Update Kerberos interface properties

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-block]#`/protocols/nfs/kerberos/interfaces/{interface.uuid}`#

Introduced In: 9.6

Updates the properties of a Kerberos interface.

== Related ONTAP commands

* `vserver nfs kerberos interface modify`

* `vserver nfs kerberos interface enable`

* `vserver nfs kerberos interface disable`

== Learn more

*

xref:{relative_path}protocols_nfs_kerberos_interfaces_endpoint_overview.html[DOC /protocols/nfs/kerberos/interfaces]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In


```

|Required
|Description

|interface.uuid
|string
|path
|True
a|Network interface UUID

* Introduced in: 9.7

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies if Kerberos is enabled.

|encryption_types
|array[string]
a|

|interface
|link:#interface[interface]
a|Network interface

|keytab_uri
|string
a|Load keytab from URI

|organizational_unit
|string
a|Organizational unit

```

```

|password
|string
a|Account creation password

|spn
|string
a|Service principal name. Valid in PATCH.

|svm
|link:#svm[svm]
a|

|user
|string
a|Account creation user name

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "encryption_types": [
    "string"
  ],
  "interface": {
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "keytab_uri": "string",
  "organizational_unit": "string",
  "password": "string",
  "spn": "string",
  "svm": {
    "name": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "user": "string"
}
====

== Response

```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response codes

|===

Error codes	Description
-------------	-------------

1966082	LIF could not be found in database. Contact technical support for assistance.
---------	-------------------------------------------------------------------------------

3276801	Failed to bind service principal name on LIF.
---------	-----------------------------------------------

3276809	Failed to disable NFS Kerberos on LIF.
---------	----------------------------------------

3276832	Failed to insert Kerberos attributes to database.
---------	---------------------------------------------------

3276842	Internal error. Failed to import Kerberos keytab file into the management databases. Contact technical support for assistance.
---------	--------------------------------------------------------------------------------------------------------------------------------

3276861	Kerberos is already enabled/disabled on this LIF.
---------	---------------------------------------------------

3276862	Kerberos service principal name is required.
---------	----------------------------------------------

3276889	Failed to enable NFS Kerberos on LIF.
---------	---------------------------------------

3276937	Failed to lookup the Vserver for the virtual interface.
---------	---------------------------------------------------------

3276941	Kerberos is a required field.
---------	-------------------------------

3276942	Service principal name is invalid. It must of the format:"nfs/+++<LIF-
---------	------------------------------------------------------------------------

```
FQDN>+++@REALM"+++</LIF-FQDN>+++
```

```
| 3276944
| Internal error. Reason: Failed to initialize the Kerberos context

| 3276945
| Internal error. Reason: Failed to parse the service principal name

| 3276951
| Warning: Skipping unsupported encryption type for service principal name

| 3276952
| "organizational_unit" option cannot be used for "Other" vendor.

| 3276965
| Account sharing across Vservers is not allowed. Use a different service
principal name unique within the first 15 characters.

| 3277019
| Cannot specify -force when enabling Kerberos.

| 3277020
| Modifying the NFS Kerberos configuration for a LIF that is not
configured for NFS is not supported.

| 3277043
| Keytab import failed due to missing keys. Keys for encryption types are
required for Vserver but found no matching keys for service principal
name. Generate the keytab file with all required keys and try again.
|===
```

```
== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```

|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#ip]
[.api-collapsible-fifth-title]
ip

IP information

[#interface]
[.api-collapsible-fifth-title]
interface

Network interface

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ip
|link:#ip[ip]
a|IP information


|name
|string
a|The name of the interface.


|uuid
|string
a|The UUID that uniquely identifies the interface.

|===

```

```

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#kerberos_interface]
[.api-collapsible-fifth-title]
kerberos_interface

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies if Kerberos is enabled.

|encryption_types
|array[string]
a|

|interface
|link:#interface[interface]
a|Network interface

```

```
|keytab_uri
|string
a|Load keytab from URI
```

```
|organizational_unit
|string
a|Organizational unit
```

```
|password
|string
a|Account creation password
```

```
|spn
|string
a|Service principal name. Valid in PATCH.
```

```
|svm
|link:#svm[svm]
a|
```

```
|user
|string
a|Account creation user name
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```



```
:leveloffset: -1
```

```
= Manage Kerberos realms
```

```
:leveloffset: +1
```

```
[[ID3541d455619601860209e933246781f1]]
```

```
= Protocols NFS Kerberos realms endpoint overview
```

```
== Examples
```

```
=== Retrieving the Kerberos realm details
```

```
----
```

```
# The API:
```

```
GET /api/protocols/nfs/kerberos/realms
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/kerberos/realms"
```

```
----
```

```
=== Creating the Kerberos realm for an SVM
```

```
----
```

```
# The API:
```

```
POST /api/protocols/nfs/kerberos/realms
```

```
# The call:
```

```
curl -d "@test_post_kerb_realm.txt" -X POST "https://<mgmt-ip>/api/protocols/nfs/kerberos/realms"
```

```
test_post_kerb_realm.txt (body):
```

```
{  
  "svm.uuid": "05c90dc2-7343-11e8-9eb4-0050568be2b7",  
  "name": "NFS-NSR-W02.RTP.NETAPP.COM",  
  "kdc": {  
    "vendor": "microsoft",  
    "ip": "10.225.185.112",  
    "port": 88  
  },  
  "comment": "realm",  
}
```

```

    "ad_server": {
        "name": "nfs-nsr-w02.rtp.netapp.com",
        "address": "10.225.185.112"
    }
}
----

=== Updating the Kerberos realm for an SVM

----

# The API:
PATCH /api/protocols/nfs/kerberos/realms/{svm.uuid}/{name}

# The call:
curl -d "@test_patch_kerb_realm.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/kerberos/realms/05c90dc2-7343-11e8-9eb4-
0050568be2b7/NFS-NSR-W02.RTP.NETAPP.COM"
test_patch_kerb_realm.txt(body):
{
    "kdc": {
        "vendor": "Microsoft",
        "ip": "100.225.185.112",
        "port": 88
    },
    "comment": "realm modify",
    "ad_server": {
        "name": "nfs.netapp.com",
        "address": "192.2.18.112"
    }
}
}
}
----

=== Deleting the Kerberos realm for an SVM

----

# The API:
DELETE /api/protocols/nfs/kerberos/realms/{svm.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/protocols/nfs/kerberos/realms/05c90dc2-7343-11e8-9eb4-
0050568be2b7/NFS-NSR-W02.RTP.NETAPP.COM"
----

```

[[ID10d926731e2861b83097f1198f342446]]

= Retrieve Kerberos realms

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nfs/kerberos/realms`#

Introduced In: 9.6

Retrieves Kerberos realms.

== Related ONTAP commands

* `vserver nfs kerberos realm show`

== Learn more

*

xref:{relative_path}protocols_nfs_kerberos_realms_endpoint_overview.html[DOC /protocols/nfs/kerberos/realms]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|ad_server.address

|string

|query

|False

a|Filter by ad_server.address

|ad_server.name

|string

```
|query  
|False  
a|Filter by ad_server.name
```

```
|svm.uuid  
|string  
|query  
|False  
a|Filter by svm.uuid
```

```
|svm.name  
|string  
|query  
|False  
a|Filter by svm.name
```

```
|kdc.vendor  
|string  
|query  
|False  
a|Filter by kdc.vendor
```

```
|kdc.port  
|integer  
|query  
|False  
a|Filter by kdc.port
```

```
|kdc.ip  
|string  
|query  
|False  
a|Filter by kdc.ip
```

```
|encryption_types  
|string  
|query  
|False  
a|Filter by encryption_types
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

```

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#kerberos_realm[kerberos_realm]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {

```

```

    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ad_server": {
        "address": "1.2.3.4",
        "name": "string"
      },
      "comment": "string",
      "encryption_types": [
        "string"
      ],
      "kdc": {
        "ip": "1.2.3.4",
        "port": "88",
        "vendor": "string"
      },
      "name": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```



```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

```

```
[#ad_server]
[.api-collapsible-fifth-title]
ad_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|Active Directory server IP address
```

```
|name
|string
a|Active Directory server name
```

```
|===
```

```
[#kdc]
[.api-collapsible-fifth-title]
kdc
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|ip
|string
a|KDC IP address
```

```
|port
|integer
a|KDC port
```

```
|vendor
|string
a|Key Distribution Center (KDC) vendor. Following values are supported:
```

```
* microsoft - Microsoft Active Directory KDC
* other - MIT Kerberos KDC or other KDC
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#kerberos_realm]
[.api-collapsible-fifth-title]
kerberos_realm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
```

```

a|

|ad_server
|link:#ad_server[ad_server]
a|

|comment
|string
a|Comment

|encryption_types
|array[string]
a|

|kdc
|link:#kdc[kdc]
a|

|name
|string
a|Kerberos realm

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message

```

```

|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDbb41bad2a5b62d5761582c5b407471f3]]
= Create a Kerberos realm

```

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/protocols/nfs/kerberos/realms`#

Introduced In: 9.6

Creates a Kerberos realm.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM on which to create the Kerberos realm.

* `name` - Base name for the Kerberos realm.

* `kdc.vendor` - Vendor of the Key Distribution Center (KDC) server for this Kerberos realm. If the configuration uses a Microsoft Active Directory domain for authentication, this field must be `microsoft`.

* `kdc.ip` - IP address of the KDC server for this Kerberos realm.

== Recommended optional properties

* `ad_server.name` - Host name of the Active Directory Domain Controller (DC). This is a mandatory parameter if the kdc-vendor is `microsoft`.

* `ad_server.address` - IP address of the Active Directory Domain Controller (DC). This is a mandatory parameter if the kdc-vendor is `microsoft`.

== Default property values

If not specified in POST, the following default property value is assigned:

* `kdc.port` - `_88_`

== Related ONTAP commands

* `vserver nfs kerberos realm create``

== Learn more

*

xref:{relative_path}protocols_nfs_kerberos_realms_endpoint_overview.html[DOC /protocols/nfs/kerberos/realms]

== Parameters

[cols=5*,options=header]

```

|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_server
|link:#ad_server[ad_server]
a|

|comment
|string
a|Comment

|encryption_types
|array[string]
a|

|kdc
|link:#kdc[kdc]
a|

|name
|string

```

```
a|Kerberos realm
```

```
|svm  
|link:#svm[svm]  
a|
```

```
|===
```

.Example request

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs+=macros]
```

```
{  
  "ad_server": {  
    "address": "1.2.3.4",  
    "name": "string"  
  },  
  "comment": "string",  
  "encryption_types": [  
    "string"  
  ],  
  "kdc": {  
    "ip": "1.2.3.4",  
    "port": "88",  
    "vendor": "string"  
  },  
  "name": "string",  
  "svm": {  
    "name": "svm1",  
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"  
  }  
}
```

```
====
```

```
== Response
```

Status: 201, Created

```
== Error
```

Status: Default

```
ONTAP Error Response codes
```



```

|===
| Error codes | Description

| 2949121
| Active Directory server name required.

| 2949122
| Active Directory server address required

| 2949123
| Failed to create Kerberos realm.

| 2949124
| Failed to create hosts file entry.

| 3276949
| Kerberos realm creation failed. Reason: The parameters "ad_server.name"
and "ad_server.address" are only valid when "kdc.vendor" is Microsoft

| 3276976
| "realm" is a required input

| 3276998
| Only the data Vservers can own NFS Kerberos realms.
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string

```

```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#ad_server]
[.api-collapsible-fifth-title]
ad_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|Active Directory server IP address

|name
|string
a|Active Directory server name

|===

[#kdc]
[.api-collapsible-fifth-title]
kdc

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ip
|string
a|KDC IP address

|port

```

```

|integer
a|KDC port

|vendor
|string
a|Key Distribution Center (KDC) vendor. Following values are supported:

* microsoft - Microsoft Active Directory KDC
* other - MIT Kerberos KDC or other KDC

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#kerberos_realm]
[.api-collapsible-fifth-title]
kerberos_realm

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|ad_server
|link:#ad_server[ad_server]
a|

|comment
|string
a|Comment

|encryption_types
|array[string]
a|

|kdc
|link:#kdc[kdc]
a|

|name
|string
a|Kerberos realm

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string

```

```

a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDff409ff4c0eca8eb4323d3e0ada9ee14]]
= Delete a Kerberos realm

```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/protocols/nfs/kerberos/realms/{svm.uuid}/{name}`#
```

Introduced In: 9.6

Deletes a Kerberos realm.

* `vserver nfs kerberos realm delete`

== Learn more

*

xref:{relative_path}protocols_nfs_kerberos_realms_endpoint_overview.html[DOC /protocols/nfs/kerberos/realms]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|SVM UUID

|name

|string

|path

|True

a|Kerberos realm

|==

== Response

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response codes

```
|===  
| Error codes | Description  
  
| 1966125  
| Failed to remove hosts entry.  
  
| 1966126  
| Failed to lookup hosts entry.  
  
| 2949141  
| Failed to lookup Kerberos realm.  
  
| 2949142  
| Failed to remove Kerberos realm.  
  
| 3276942  
| Service principal name is invalid. It must of the format:"nfs/+++<LIF-  
FQDN>+++@REALM\\\\"+++</LIF-FQDN>+++  
  
| 3276976  
| "realm" is a required input  
  
| 3276998  
| Only the data Vservers can own NFS Kerberos realms.  
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument
```



```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDc8150c4901aadcd6fa2e6378154e9ac7]]
= Retrieve a Kerberos realm

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/kerberos/realms/{svm.uuid}/{name}`#

```

Introduced In: 9.6

Retrieves a Kerberos realm.

* `vserver nfs kerberos realm show`

== Learn more

*

xref:{relative_path}protocols_nfs_kerberos_realms_endpoint_overview.html[D
OC /protocols/nfs/kerberos/realms]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|SVM UUID

|name

|string

|path

|True

a|Kerberos realm

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ad_server
|link:#ad_server[ad_server]
a|

|comment
|string
a|Comment

|encryption_types
|array[string]
a|

|kdc
|link:#kdc[kdc]
a|

|name
|string
a|Kerberos realm

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
=====
```

```
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_server": {
    "address": "1.2.3.4",
    "name": "string"
  },
  "comment": "string",
  "encryption_types": [
    "string"
  ],
  "kdc": {
    "ip": "1.2.3.4",
    "port": "88",
    "vendor": "string"
  },
  "name": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
```

```

a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#ad_server]
[.api-collapsible-fifth-title]
ad_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|Active Directory server IP address

|name
|string
a|Active Directory server name

|===

[#kdc]
[.api-collapsible-fifth-title]
kdc

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|ip
|string
a|KDC IP address


|port
|integer
a|KDC port


|vendor
|string
a|Key Distribution Center (KDC) vendor. Following values are supported:

* microsoft - Microsoft Active Directory KDC
* other - MIT Kerberos KDC or other KDC


|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.


|uuid
|string
a|The unique identifier of the SVM.

```

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code


```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID4e3af3f16b70c13401da09037b515fc7]]
= Update Kerberos realm properties
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/nfs/kerberos/realms/{svm.uuid}/{name}`#
```

Introduced In: 9.6

Updates the properties of a Kerberos realm.

* `vserver nfs kerberos realm modify`

== Learn more

*

xref:{relative_path}protocols_nfs_kerberos_realms_endpoint_overview.html[D
OC /protocols/nfs/kerberos/realms]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
```

|Description

|svm.uuid

|string

|path

|True

a|SVM UUID

|name

|string

|path

|True

a|Kerberos realm

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|ad_server

|link:#ad_server[ad_server]

a|

|comment

|string

a|Comment

|encryption_types

|array[string]

a|

|kdc

|link:#kdc[kdc]

a|

|name

|string

a|Kerberos realm

```

|svm
|link:#svm[svm]
a|

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "ad_server": {
    "address": "1.2.3.4",
    "name": "string"
  },
  "comment": "string",
  "encryption_types": [
    "string"
  ],
  "kdc": {
    "ip": "1.2.3.4",
    "port": "88",
    "vendor": "string"
  },
  "name": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

== Response

```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response codes

```

|===
| Error codes | Description

| 1966125
| Failed to remove hosts entry.

| 1966126
| Failed to lookup hosts entry.

| 1966131
| Failed to create hosts entry.

| 1966132
| Failed to modify hosts entry.

| 2949121
| Active Directory server name required.

| 2949122
| Active Directory server address required

| 2949123
| Failed to create Kerberos realm.

| 2949124
| Failed to create hosts file entry.

| 2949141
| Failed to lookup Kerberos realm.

| 2949148
| Failed to modify Kerberos realm.

| 3276976
| "realm" is a required input

| 3276998
| Only the data Vservers can own NFS Kerberos realms.
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#ad_server]
[.api-collapsible-fifth-title]
ad_server

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|Active Directory server IP address

|name
|string
a|Active Directory server name

|===

[#kdc]
[.api-collapsible-fifth-title]
kdc

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ip
|string
a|KDC IP address

|port
|integer
a|KDC port

|vendor
|string
a|Key Distribution Center (KDC) vendor. Following values are supported:

* microsoft - Microsoft Active Directory KDC
* other - MIT Kerberos KDC or other KDC

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

```

|===

[#kerberos_realm]
[.api-collapsible-fifth-title]
kerberos_realm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_server
|link:#ad_server[ad_server]
a|

|comment
|string
a|Comment

|encryption_types
|array[string]
a|

|kdc
|link:#kdc[kdc]
a|

|name
|string
a|Kerberos realm

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```



```
|===
```

```
//end collapsible .Definitions block  
====
```

```
:leveloffset: -1
```

```
= Manage NFS services
```

```
:leveloffset: +1
```

```
[[ID8b0cb314208d5b9cbc43fe3f65207565]]  
= Protocols NFS services endpoint overview
```

```
=== Retrieving an NFS configuration
```

```
----
```

```
# The API:  
GET /api/protocols/nfs/services
```

```
# The call:  
curl -X GET "https://<mgmt-ip>/api/protocols/nfs/services"  
----
```

```
=== Creating an NFS configuration for an SVM
```

```
----
```

```
# The API:  
POST /api/protocols/nfs/services
```

```
# The call:  
curl -d "@test_nfs_post.txt" -X POST "https://<mgmt-  
ip>/api/protocols/nfs/services"  
test_nfs_post.txt (body):  
{  
  "svm": {
```

```

    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "protocol": {
    "v4_id_domain": "nfs-nsr-w01.rtp.netapp.com"
  },
  "vstorage_enabled": "true"
}
-----

```

=== Updating an NFS configuration for an SVM

```

-----

# The API:
PATCH /api/protocols/nfs/services/{svm.uuid}

# The call:
curl -d "@test_nfs_patch.txt" -X PATCH "https://<mgmt-
ip>/api/protocols/nfs/services/4a415601-548c-11e8-a21d-0050568bcb9"
test_nfs_patch.txt(body):
{
  "protocol": {
    "v4_id_domain": "nfs-nsr-w01.rtp.netapp.com"
  },
  "vstorage_enabled": "false"
}
-----

```

=== Deleting an NFS configuration for an SVM

```

-----

# The API:
DELETE /api/protocols/nfs/services/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/nfs/services/4a415601-
548c-11e8-a21d-0050568bcb9"
-----

```

== Performance monitoring

Performance of the SVM can be monitored by the ``metric.++`` and ``statistics.++`` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The ``metric.++`` properties denote an average whereas ``statistics.++`` properties denote a real-time monotonically increasing value aggregated across all nodes.

[[ID9f2e5ee8928d850229667e861edd4bdb]]

= Retrieve the NFS configuration for SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nfs/services`#

Introduced In: 9.6

Retrieves the NFS configuration of SVMs.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `statistics.+++`

* `metric.+++`

== Related ONTAP commands

* `vserver nfs show`

* `vserver nfs status`

== Learn more

* [xref:{relative_path}protocols_nfs_services_endpoint_overview.html\[DOC /protocols/nfs/services\]](#)

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|transport.udp_enabled
|boolean
|query
|False
a|Filter by transport.udp_enabled
```

```
|transport.tcp_enabled
|boolean
|query
|False
a|Filter by transport.tcp_enabled
```

```
|extended_groups_limit
|integer
|query
|False
a|Filter by extended_groups_limit
```

* Introduced in: 9.8

```
|positive_cached_credential_ttl
|integer
|query
|False
a|Filter by positive_cached_credential_ttl
```

* Introduced in: 9.8

```
|metric.v3.throughput.write
|integer
|query
|False
a|Filter by metric.v3.throughput.write
```

* Introduced in: 9.7

```
|metric.v3.throughput.read
|integer
|query
|False
a|Filter by metric.v3.throughput.read

* Introduced in: 9.7

|metric.v3.throughput.total
|integer
|query
|False
a|Filter by metric.v3.throughput.total

* Introduced in: 9.7

|metric.v3.duration
|string
|query
|False
a|Filter by metric.v3.duration

* Introduced in: 9.7

|metric.v3.timestamp
|string
|query
|False
a|Filter by metric.v3.timestamp

* Introduced in: 9.7

|metric.v3.status
|string
|query
|False
a|Filter by metric.v3.status

* Introduced in: 9.7
```

```
|metric.v3.iops.total
|integer
|query
|False
a|Filter by metric.v3.iops.total
```

* Introduced in: 9.7

```
|metric.v3.iops.read
|integer
|query
|False
a|Filter by metric.v3.iops.read
```

* Introduced in: 9.7

```
|metric.v3.iops.other
|integer
|query
|False
a|Filter by metric.v3.iops.other
```

* Introduced in: 9.7

```
|metric.v3.iops.write
|integer
|query
|False
a|Filter by metric.v3.iops.write
```

* Introduced in: 9.7

```
|metric.v3.latency.total
|integer
|query
|False
a|Filter by metric.v3.latency.total
```

* Introduced in: 9.7

```
|metric.v3.latency.read
|integer
```

```
|query
|False
a|Filter by metric.v3.latency.read
```

* Introduced in: 9.7

```
|metric.v3.latency.other
|integer
|query
|False
a|Filter by metric.v3.latency.other
```

* Introduced in: 9.7

```
|metric.v3.latency.write
|integer
|query
|False
a|Filter by metric.v3.latency.write
```

* Introduced in: 9.7

```
|metric.v41.throughput.write
|integer
|query
|False
a|Filter by metric.v41.throughput.write
```

* Introduced in: 9.8

```
|metric.v41.throughput.read
|integer
|query
|False
a|Filter by metric.v41.throughput.read
```

* Introduced in: 9.8

```
|metric.v41.throughput.total
|integer
|query
|False
```

a|Filter by metric.v41.throughput.total

* Introduced in: 9.8

|metric.v41.duration

|string

|query

|False

a|Filter by metric.v41.duration

* Introduced in: 9.8

|metric.v41.timestamp

|string

|query

|False

a|Filter by metric.v41.timestamp

* Introduced in: 9.8

|metric.v41.status

|string

|query

|False

a|Filter by metric.v41.status

* Introduced in: 9.8

|metric.v41.iops.total

|integer

|query

|False

a|Filter by metric.v41.iops.total

* Introduced in: 9.8

|metric.v41.iops.read

|integer

|query

|False

a|Filter by metric.v41.iops.read

* Introduced in: 9.8

```
|metric.v41.iops.other  
|integer  
|query  
|False  
a|Filter by metric.v41.iops.other
```

* Introduced in: 9.8

```
|metric.v41.iops.write  
|integer  
|query  
|False  
a|Filter by metric.v41.iops.write
```

* Introduced in: 9.8

```
|metric.v41.latency.total  
|integer  
|query  
|False  
a|Filter by metric.v41.latency.total
```

* Introduced in: 9.8

```
|metric.v41.latency.read  
|integer  
|query  
|False  
a|Filter by metric.v41.latency.read
```

* Introduced in: 9.8

```
|metric.v41.latency.other  
|integer  
|query  
|False  
a|Filter by metric.v41.latency.other
```

* Introduced in: 9.8

```
|metric.v41.latency.write
|integer
|query
|False
a|Filter by metric.v41.latency.write
```

* Introduced in: 9.8

```
|metric.v4.throughput.write
|integer
|query
|False
a|Filter by metric.v4.throughput.write
```

* Introduced in: 9.8

```
|metric.v4.throughput.read
|integer
|query
|False
a|Filter by metric.v4.throughput.read
```

* Introduced in: 9.8

```
|metric.v4.throughput.total
|integer
|query
|False
a|Filter by metric.v4.throughput.total
```

* Introduced in: 9.8

```
|metric.v4.duration
|string
|query
|False
a|Filter by metric.v4.duration
```

* Introduced in: 9.8

```
|metric.v4.timestamp
```

```
|string
|query
|False
a|Filter by metric.v4.timestamp
```

* Introduced in: 9.8

```
|metric.v4.status
|string
|query
|False
a|Filter by metric.v4.status
```

* Introduced in: 9.8

```
|metric.v4.iops.total
|integer
|query
|False
a|Filter by metric.v4.iops.total
```

* Introduced in: 9.8

```
|metric.v4.iops.read
|integer
|query
|False
a|Filter by metric.v4.iops.read
```

* Introduced in: 9.8

```
|metric.v4.iops.other
|integer
|query
|False
a|Filter by metric.v4.iops.other
```

* Introduced in: 9.8

```
|metric.v4.iops.write
|integer
|query
```

```
|False
a|Filter by metric.v4.iops.write

* Introduced in: 9.8

|metric.v4.latency.total
|integer
|query
|False
a|Filter by metric.v4.latency.total

* Introduced in: 9.8

|metric.v4.latency.read
|integer
|query
|False
a|Filter by metric.v4.latency.read

* Introduced in: 9.8

|metric.v4.latency.other
|integer
|query
|False
a|Filter by metric.v4.latency.other

* Introduced in: 9.8

|metric.v4.latency.write
|integer
|query
|False
a|Filter by metric.v4.latency.write

* Introduced in: 9.8

|protocol.v41_enabled
|boolean
|query
|False
a|Filter by protocol.v41_enabled
```

```
|protocol.v3_enabled
|boolean
|query
|False
a|Filter by protocol.v3_enabled
```

```
|protocol.v41_features.pnfs_enabled
|boolean
|query
|False
a|Filter by protocol.v41_features.pnfs_enabled
```

```
|protocol.v41_features.write_delegation_enabled
|boolean
|query
|False
a|Filter by protocol.v41_features.write_delegation_enabled
```

```
|protocol.v41_features.acl_enabled
|boolean
|query
|False
a|Filter by protocol.v41_features.acl_enabled
```

```
|protocol.v41_features.read_delegation_enabled
|boolean
|query
|False
a|Filter by protocol.v41_features.read_delegation_enabled
```

```
|protocol.v40_features.read_delegation_enabled
|boolean
|query
|False
a|Filter by protocol.v40_features.read_delegation_enabled
```

```
|protocol.v40_features.write_delegation_enabled
|boolean
|query
```

```
|False
a|Filter by protocol.v40_features.write_delegation_enabled
```

```
|protocol.v40_features.acl_enabled
|boolean
|query
|False
a|Filter by protocol.v40_features.acl_enabled
```

```
|protocol.v4_64bit_identifiers_enabled
|boolean
|query
|False
a|Filter by protocol.v4_64bit_identifiers_enabled
```

* Introduced in: 9.8

```
|protocol.v3_64bit_identifiers_enabled
|boolean
|query
|False
a|Filter by protocol.v3_64bit_identifiers_enabled
```

* Introduced in: 9.8

```
|protocol.v40_enabled
|boolean
|query
|False
a|Filter by protocol.v40_enabled
```

```
|protocol.v4_id_domain
|string
|query
|False
a|Filter by protocol.v4_id_domain
```

```
|statistics.v41.latency_raw.total
|integer
|query
|False
```

```
a|Filter by statistics.v41.latency_raw.total
```

```
* Introduced in: 9.8
```

```
|statistics.v41.latency_raw.read
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by statistics.v41.latency_raw.read
```

```
* Introduced in: 9.8
```

```
|statistics.v41.latency_raw.other
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by statistics.v41.latency_raw.other
```

```
* Introduced in: 9.8
```

```
|statistics.v41.latency_raw.write
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by statistics.v41.latency_raw.write
```

```
* Introduced in: 9.8
```

```
|statistics.v41.timestamp
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by statistics.v41.timestamp
```

```
* Introduced in: 9.8
```

```
|statistics.v41.iops_raw.total
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by statistics.v41.iops_raw.total
```

* Introduced in: 9.8

```
|statistics.v41.iops_raw.read  
|integer  
|query  
|False  
a|Filter by statistics.v41.iops_raw.read
```

* Introduced in: 9.8

```
|statistics.v41.iops_raw.other  
|integer  
|query  
|False  
a|Filter by statistics.v41.iops_raw.other
```

* Introduced in: 9.8

```
|statistics.v41.iops_raw.write  
|integer  
|query  
|False  
a|Filter by statistics.v41.iops_raw.write
```

* Introduced in: 9.8

```
|statistics.v41.throughput_raw.write  
|integer  
|query  
|False  
a|Filter by statistics.v41.throughput_raw.write
```

* Introduced in: 9.8

```
|statistics.v41.throughput_raw.read  
|integer  
|query  
|False  
a|Filter by statistics.v41.throughput_raw.read
```

* Introduced in: 9.8


```
|statistics.v41.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.v41.throughput_raw.total
```

* Introduced in: 9.8

```
|statistics.v41.status
|string
|query
|False
a|Filter by statistics.v41.status
```

* Introduced in: 9.8

```
|statistics.v3.latency_raw.total
|integer
|query
|False
a|Filter by statistics.v3.latency_raw.total
```

* Introduced in: 9.7

```
|statistics.v3.latency_raw.read
|integer
|query
|False
a|Filter by statistics.v3.latency_raw.read
```

* Introduced in: 9.7

```
|statistics.v3.latency_raw.other
|integer
|query
|False
a|Filter by statistics.v3.latency_raw.other
```

* Introduced in: 9.7

```
|statistics.v3.latency_raw.write
```

```
|integer
|query
|False
a|Filter by statistics.v3.latency_raw.write
```

* Introduced in: 9.7

```
|statistics.v3.timestamp
|string
|query
|False
a|Filter by statistics.v3.timestamp
```

* Introduced in: 9.7

```
|statistics.v3.iops_raw.total
|integer
|query
|False
a|Filter by statistics.v3.iops_raw.total
```

* Introduced in: 9.7

```
|statistics.v3.iops_raw.read
|integer
|query
|False
a|Filter by statistics.v3.iops_raw.read
```

* Introduced in: 9.7

```
|statistics.v3.iops_raw.other
|integer
|query
|False
a|Filter by statistics.v3.iops_raw.other
```

* Introduced in: 9.7

```
|statistics.v3.iops_raw.write
|integer
|query
```

```

|False
a|Filter by statistics.v3.iops_raw.write

* Introduced in: 9.7

|statistics.v3.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.v3.throughput_raw.write

* Introduced in: 9.7

|statistics.v3.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.v3.throughput_raw.read

* Introduced in: 9.7

|statistics.v3.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.v3.throughput_raw.total

* Introduced in: 9.7

|statistics.v3.status
|string
|query
|False
a|Filter by statistics.v3.status

* Introduced in: 9.7

|statistics.v4.latency_raw.total
|integer
|query
|False
a|Filter by statistics.v4.latency_raw.total

```

* Introduced in: 9.8

```
|statistics.v4.latency_raw.read  
|integer  
|query  
|False  
a|Filter by statistics.v4.latency_raw.read
```

* Introduced in: 9.8

```
|statistics.v4.latency_raw.other  
|integer  
|query  
|False  
a|Filter by statistics.v4.latency_raw.other
```

* Introduced in: 9.8

```
|statistics.v4.latency_raw.write  
|integer  
|query  
|False  
a|Filter by statistics.v4.latency_raw.write
```

* Introduced in: 9.8

```
|statistics.v4.timestamp  
|string  
|query  
|False  
a|Filter by statistics.v4.timestamp
```

* Introduced in: 9.8

```
|statistics.v4.iops_raw.total  
|integer  
|query  
|False  
a|Filter by statistics.v4.iops_raw.total
```

* Introduced in: 9.8

```
|statistics.v4.iops_raw.read
|integer
|query
|False
a|Filter by statistics.v4.iops_raw.read
```

* Introduced in: 9.8

```
|statistics.v4.iops_raw.other
|integer
|query
|False
a|Filter by statistics.v4.iops_raw.other
```

* Introduced in: 9.8

```
|statistics.v4.iops_raw.write
|integer
|query
|False
a|Filter by statistics.v4.iops_raw.write
```

* Introduced in: 9.8

```
|statistics.v4.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.v4.throughput_raw.write
```

* Introduced in: 9.8

```
|statistics.v4.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.v4.throughput_raw.read
```

* Introduced in: 9.8

```
|statistics.v4.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.v4.throughput_raw.total
```

* Introduced in: 9.8

```
|statistics.v4.status
|string
|query
|False
a|Filter by statistics.v4.status
```

* Introduced in: 9.8

```
|state
|string
|query
|False
a|Filter by state
```

```
|auth_sys_extended_groups_enabled
|boolean
|query
|False
a|Filter by auth_sys_extended_groups_enabled
```

* Introduced in: 9.8

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|rquota_enabled
|boolean
|query
|False
a|Filter by rquota_enabled
```

* Introduced in: 9.8

```
|showmount_enabled
|boolean
|query
|False
a|Filter by showmount_enabled
```

* Introduced in: 9.8

```
|vstorage_enabled
|boolean
|query
|False
a|Filter by vstorage_enabled
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
```

reached.

* Default value: 1

* Max value: 120

* Min value: 0

|return_records

|boolean

|query

|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of NFS Server Records

|records


```
|array[link:#nfs_service[nfs_service]]
```

```
a|
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "extended_groups_limit": "32",
      "metric": {
        "v3": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "duration": "PT15S",
          "iops": {
            "read": "200",
            "total": "1000",
            "write": "100"
          },
          "latency": {
            "read": "200",
            "total": "1000",
            "write": "100"
          },
          "status": "ok",

```

```

    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "v4": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "v41": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",

```

```

        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
}
},
"positive_cached_credential_ttl": "7200000",
"protocol": {
    "v4_id_domain": "string"
},
"state": "string",
"statistics": {
    "v3": {
        "iops_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "latency_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "status": "ok",
        "throughput_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "timestamp": "2017-01-25T11:20:13Z"
    },
    "v4": {
        "iops_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "latency_raw": {
            "read": "200",
            "total": "1000",

```

```

        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"v41": {
    "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
}
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
]
}
=====

```

== Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

```

|===

[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput]
```

```
[.api-collapsible-fifth-title]
```

throughput

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.


```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#v3]
[.api-collapsible-fifth-title]
v3
```

The NFSv3 operations

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#v4]
[.api-collapsible-fifth-title]
v4
```

The NFSv4 operations

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|duration
```

```

|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#v41]
[.api-collapsible-fifth-title]
v41

```

The NFSv4.1 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|duration
```

```
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
```

```
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
```

```
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
```

```

|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

Historical performance numbers, such as IOPS latency and throughput, for
SVM-NFS protocol.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|v3
|link:#v3[v3]
a|The NFSv3 operations

|v4
|link:#v4[v4]
a|The NFSv4 operations

|v41
|link:#v41[v41]
a|The NFSv4.1 operations

|===

[#v40_features]
[.api-collapsible-fifth-title]
v40_features

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|acl_enabled
|boolean
a|Specifies whether NFSv4.0 ACLs is enabled.

|read_delegation_enabled
|boolean
a|Specifies whether NFSv4.0 Read Delegation is enabled.

|write_delegation_enabled
|boolean
a|Specifies whether NFSv4.0 Write Delegation is enabled.

|===

[#v41_features]
[.api-collapsible-fifth-title]
v41_features

[cols=3*,options=header]
|===
|Name
|Type
|Description

|acl_enabled
|boolean
a|Specifies whether NFSv4.1 or later ACLs is enabled.

|pnfs_enabled
|boolean
a|Specifies whether NFSv4.1 or later Parallel NFS is enabled.

|read_delegation_enabled
|boolean

```

a|Specifies whether NFSv4.1 or later Read Delegation is enabled.

|write_delegation_enabled

|boolean

a|Specifies whether NFSv4.1 or later Write Delegation is enabled.

|===

[#protocol]

[.api-collapsible-fifth-title]

protocol

[cols=3*,options=header]

|===

|Name

|Type

|Description

|v3_64bit_identifiers_enabled

|boolean

a|Specifies whether 64-bit support for NFSv3 FSIDs and file IDs is enabled.

|v3_enabled

|boolean

a|Specifies whether NFSv3 protocol is enabled.

|v40_enabled

|boolean

a|Specifies whether NFSv4.0 protocol is enabled.

|v40_features

|link:#v40_features[v40_features]

a|

|v41_enabled

|boolean

a|Specifies whether NFSv4.1 or later protocol is enabled.

|v41_features

```
|link:#v41_features[v41_features]
```

```
a|
```

```
|v4_64bit_identifiers_enabled
```

```
|boolean
```

```
a|Specifies whether 64-bit support for NFSv4.x FSIDs and file IDs is enabled.
```

```
|v4_id_domain
```

```
|string
```

```
a|Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol.
```

```
|===
```

```
[#iops_raw]
```

```
[.api-collapsible-fifth-title]
```

```
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```



```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#v3]
[.api-collapsible-fifth-title]
v3
```

The NFSv3 operations

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#v4]
[.api-collapsible-fifth-title]

v4

The NFSv4 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
```

```
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#v41]
[.api-collapsible-fifth-title]
v41
```

The NFSv4.1 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not

have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|v3
|link:#v3[v3]
a|The NFSv3 operations
```

```
|v4
|link:#v4[v4]
a|The NFSv4 operations
```

```
|v41
|link:#v41[v41]
a|The NFSv4.1 operations
```

```

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#transport]
[.api-collapsible-fifth-title]
transport

[cols=3*,options=header]
|===
|Name
|Type
|Description

|tcp_enabled
|boolean
a|Specifies whether TCP transports are enabled on the server.

|udp_enabled
|boolean

```

a|Specifies whether UDP transports are enabled on the server.

|===

[#nfs_service]

[.api-collapsible-fifth-title]

nfs_service

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|auth_sys_extended_groups_enabled

|boolean

a|Specifies whether or not extended groups support over AUTH_SYS is enabled.

|enabled

|boolean

a|Specifies if the NFS service is administratively enabled.

|extended_groups_limit

|integer

a|Specifies the maximum auxillary groups supported over AUTH_SYS and RPCSEC_GSS.

|metric

|link:#metric[metric]

a|Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

|positive_cached_credential_ttl

|integer

a|Specifies the time to live value (in msecs) of a positive cached credential


```

|protocol
|link:#protocol[protocol]
a|

|rquota_enabled
|boolean
a|Specifies whether or not the remote quota feature is enabled.

|showmount_enabled
|boolean
a|Specifies whether or not the showmount feature is enabled.

|state
|string
a|Specifies the state of the NFS service on the SVM. The following values
are supported:
    * online - NFS server is ready to accept client requests.
    * offline - NFS server is not ready to accept client requests.

|statistics
|link:#statistics[statistics]
a|Realtime performance numbers, such as IOPS latency and throughput, for
SVM-NFS protocol.

|svm
|link:#svm[svm]
a|

|transport
|link:#transport[transport]
a|

|vstorage_enabled
|boolean
a|Specifies whether or not the VMware vstorage feature is enabled.

|===

[#error_arguments]

```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

```

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID6c6f802aa2b0eb308b89e4feae0437ce]]
= Create an NFS configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/nfs/services`#

*Introduced In:* 9.6

Creates an NFS configuration for an SVM.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM for which to create the NFS
configuration.

== Default property values

If not specified in POST, the following default property values are
assigned:

* `enabled` - _true_
* `state` - online
* `transport.udp_enabled` - _true_
* `transport.tcp_enabled` - _true_
* `protocol.v3_enabled` - _true_
* `protocol.v3_64bit_identifiers_enabled` - _false_
* `protocol.v4_id_domain` - defaultv4iddomain.com
* `protocol.v4_64bit_identifiers_enabled` - _true_
* `protocol.v4_enabled` - _false_
* `protocol.v41_enabled` - _false_
* `protocol.v40_features.acl_enabled` - _false_
* `protocol.v40_features.read_delegation_enabled` - _false_
* `protocol.v40_features.write_delegation_enabled` - _false_
* `protocol.v41_features.acl_enabled` - _false_

```

```

* `protocol.v41_features.read_delegation_enabled` - __false__
* `protocol.v41_features.write_delegation_enabled` - __false__
* `protocol.v41_features.pnfs_enabled` - __false__
* `vstorage_enabled` - __false__
* `rquota_enabled` - __false__
* `showmount_enabled` - __true__
* `auth_sys_extended_groups_enabled` - __false__
* `extended_groups_limit` - __32__
* `positive_cached_credential_ttl` - __7200000__

```

== Related ONTAP commands

```

* `vserver nfs create`

```

== Learn more

```

* xref:{relative_path}protocols_nfs_services_endpoint_overview.html[DOC
/protocols/nfs/services]

```

== Parameters

```

[cols=5*,options=header]
|===

```

```

|Name
|Type
|In
|Required
|Description

```

```

|return_records
|boolean
|query
|False

```

```

a|The default is false. If set to true, the records are returned.

```

```

* Default value:

```

```

|===

```

== Request Body

```

[cols=3*,options=header]
|===

```

Name	Type	Description
auth_sys_extended_groups_enabled	boolean	a Specifies whether or not extended groups support over AUTH_SYS is enabled.
enabled	boolean	a Specifies if the NFS service is administratively enabled.
extended_groups_limit	integer	a Specifies the maximum auxillary groups supported over AUTH_SYS and RPCSEC_GSS.
positive_cached_credential_ttl	integer	a Specifies the time to live value (in msec) of a positive cached credential
protocol	link:#protocol[protocol]	a
rquota_enabled	boolean	a Specifies whether or not the remote quota feature is enabled.
showmount_enabled	boolean	a Specifies whether or not the showmount feature is enabled.
state	string	a Specifies the state of the NFS service on the SVM. The following values are supported: <ul style="list-style-type: none"> * online - NFS server is ready to accept client requests. * offline - NFS server is not ready to accept client requests.

```

|svm
|link:#svm[svm]
a|

|transport
|link:#transport[transport]
a|

|vstorage_enabled
|boolean
a|Specifies whether or not the VMware vstorage feature is enabled.

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "extended_groups_limit": "32",
  "positive_cached_credential_ttl": "7200000",
  "protocol": {
    "v4_id_domain": "string"
  },
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response

```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of NFS Server Records


|records
|array[link:#nfs_service[nfs_service]]
a|

|===


.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "extended_groups_limit": "32",
      "positive_cached_credential_ttl": "7200000",
      "protocol": {
        "v4_id_domain": "string"
      },
      "state": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 3276916
| Vserver is not running

| 3276994
| Kerberos must be disabled on all LIFs on Vserver before adding or
removing AES encryption. Disable Kerberos on the LIF and try again

| 3277038
| Cannot enable \"showmount\" feature because it requires an effective
cluster version of Data ONTAP 8.3.0 or later

| 3277049
| Cannot enable \"showmount\" feature on ID-Discard Vserver. Ensure that
the Vserver is initialized and retry the command

| 3277052
| NFSv4.x access to transitioned volumes in this Vserver could trigger
conversion of non-Unicode directories to Unicode, which might impact data-
serving performance. Before enabling NFSv4.x for this Vserver, refer to
the Data and Configuration Transition Guide

| 3277069
| Cannot disable TCP because the SnapDiff RPC server is in the \"on\"
state

| 3277089
| Attempting to create an NFS server using 64-bits for NFSv3 FSIDs and
File IDs on Vserver. Older client software might not work with 64-bit
identifiers

| 3277099
| Domain name contains invalid characters or it is too short. Allowed
characters are: alphabetical characters (A-Za-z), numeric characters (0-
9), minus sign (-), and the period (.). The first character must be
alphabetical or numeric, last character must not be a minus sign or a
period. Minimum supported length: 2 characters, maximum of 256 characters
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions

```



```
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#latency]
[.api-collapsible-fifth-title]
latency

```

The round trip latency in microseconds observed at the storage object.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

```

```
|===
```

```
[#throughput]  
[.api-collapsible-fifth-title]  
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|read  
|integer  
a|Performance metric for read I/O operations.
```

```
|total  
|integer  
a|Performance metric aggregated over all types of I/O operations.
```

```
|write  
|integer  
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#v3]  
[.api-collapsible-fifth-title]  
v3
```

The NFSv3 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#v4]
[.api-collapsible-fifth-title]

```

v4

The NFSv4 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|duration
```

```
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
```

```
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
```

```
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
```

```
|link:#throughput[throughput]
```

a|The rate of throughput bytes per second observed at the storage object.

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#v41]
[.api-collapsible-fifth-title]
v41
```

The NFSv4.1 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

```
|status
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time

between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#metric]

[.api-collapsible-fifth-title]

metric

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|v3

|link:#v3[v3]

a|The NFSv3 operations

|v4

|link:#v4[v4]

a|The NFSv4 operations

|v41

|link:#v41[v41]

a|The NFSv4.1 operations

|===

[#v40_features]

[.api-collapsible-fifth-title]

v40_features

[cols=3*,options=header]

|===

|Name

|Type

|Description

|acl_enabled

|boolean

a|Specifies whether NFSv4.0 ACLs is enabled.

|read_delegation_enabled

|boolean

a|Specifies whether NFSv4.0 Read Delegation is enabled.

|write_delegation_enabled

|boolean

a|Specifies whether NFSv4.0 Write Delegation is enabled.

|===

[#v41_features]

[.api-collapsible-fifth-title]

v41_features

[cols=3*,options=header]

|===

|Name

|Type

|Description

|acl_enabled

|boolean

a|Specifies whether NFSv4.1 or later ACLs is enabled.

|pnfs_enabled
|boolean
a|Specifies whether NFSv4.1 or later Parallel NFS is enabled.

|read_delegation_enabled
|boolean
a|Specifies whether NFSv4.1 or later Read Delegation is enabled.

|write_delegation_enabled
|boolean
a|Specifies whether NFSv4.1 or later Write Delegation is enabled.

|===

[#protocol]
[.api-collapsible-fifth-title]
protocol

[cols=3*,options=header]
|===
|Name
|Type
|Description

|v3_64bit_identifiers_enabled
|boolean
a|Specifies whether 64-bit support for NFSv3 FSIDs and file IDs is enabled.

|v3_enabled
|boolean
a|Specifies whether NFSv3 protocol is enabled.

|v40_enabled
|boolean
a|Specifies whether NFSv4.0 protocol is enabled.

|v40_features
|link:#v40_features[v40_features]

a|

|v41_enabled

|boolean

a|Specifies whether NFSv4.1 or later protocol is enabled.

|v41_features

|link:#v41_features[v41_features]

a|

|v4_64bit_identifiers_enabled

|boolean

a|Specifies whether 64-bit support for NFSv4.x FSIDs and file IDs is enabled.

|v4_id_domain

|string

a|Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol.

|===

[#iops_raw]

[.api-collapsible-fifth-title]

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#v3]
[.api-collapsible-fifth-title]
v3
```

The NFSv3 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
```

```
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp
```

```
|string
```

a|The timestamp of the performance data.

|===

[#v4]

[.api-collapsible-fifth-title]

v4

The NFSv4 operations

[cols=3*,options=header]

|===

|Name

|Type

|Description

|iops_raw

|link:#iops_raw[iops_raw]

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw

|link:#latency_raw[latency_raw]

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#v41]
[.api-collapsible-fifth-title]
v41
```

The NFSv4.1 operations

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.
```

```
|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
```

returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|v3
|link:#v3[v3]
a|The NFSv3 operations
```

```
|v4
|link:#v4[v4]
```


a|The NFSv4 operations

|v41

|link:#v41[v41]

a|The NFSv4.1 operations

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#transport]

[.api-collapsible-fifth-title]

transport

[cols=3*,options=header]

|===

|Name

|Type

|Description

|tcp_enabled

|boolean

a|Specifies whether TCP transports are enabled on the server.

```

|udp_enabled
|boolean
a|Specifies whether UDP transports are enabled on the server.

|===

[#nfs_service]
[.api-collapsible-fifth-title]
nfs_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|auth_sys_extended_groups_enabled
|boolean
a|Specifies whether or not extended groups support over AUTH_SYS is
enabled.

|enabled
|boolean
a|Specifies if the NFS service is administratively enabled.

|extended_groups_limit
|integer
a|Specifies the maximum auxillary groups supported over AUTH_SYS and
RPCSEC_GSS.

|positive_cached_credential_ttl
|integer
a|Specifies the time to live value (in msec) of a positive cached
credential

|protocol
|link:#protocol[protocol]
a|

```

```

|rquota_enabled
|boolean
a|Specifies whether or not the remote quota feature is enabled.

|showmount_enabled
|boolean
a|Specifies whether or not the showmount feature is enabled.

|state
|string
a|Specifies the state of the NFS service on the SVM. The following values
are supported:
    * online - NFS server is ready to accept client requests.
    * offline - NFS server is not ready to accept client requests.

|svm
|link:#svm[svm]
a|

|transport
|link:#transport[transport]
a|

|vstorage_enabled
|boolean
a|Specifies whether or not the VMware vstorage feature is enabled.

|===

[#nfs_service]
[.api-collapsible-fifth-title]
nfs_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|auth_sys_extended_groups_enabled
|boolean
a|Specifies whether or not extended groups support over AUTH_SYS is

```

enabled.

|enabled

|boolean

a|Specifies if the NFS service is administratively enabled.

|extended_groups_limit

|integer

a|Specifies the maximum auxillary groups supported over AUTH_SYS and RPCSEC_GSS.

|positive_cached_credential_ttl

|integer

a|Specifies the time to live value (in msec) of a positive cached credential

|protocol

|link:#protocol[protocol]

a|

|rquota_enabled

|boolean

a|Specifies whether or not the remote quota feature is enabled.

|showmount_enabled

|boolean

a|Specifies whether or not the showmount feature is enabled.

|state

|string

a|Specifies the state of the NFS service on the SVM. The following values are supported:

***** online - NFS server is ready to accept client requests.

***** offline - NFS server is not ready to accept client requests.

|svm

|link:#svm[svm]

a|

```
|transport
|link:#transport[transport]
a|
```

```
|vstorage_enabled
|boolean
a|Specifies whether or not the VMware vstorage feature is enabled.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[IDbc06020e15476ed82300f95107fd2cf3]]
```

```
= Delete the NFS configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/protocols/nfs/services/{svm.uuid}`#
```

```
*Introduced In:* 9.6
```

```
Deletes the NFS configuration of an SVM.
```

```
== Related ONTAP commands
```

```
* `vserver nfs delete`
```

```
== Learn more
```

```
* xref:{relative_path}protocols_nfs_services_endpoint_overview.html[DOC  
/protocols/nfs/services]
```

```

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|
|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

```

ONTAP Error Response Codes

|===
| Error Code | Description

| 3276916
| Vserver is not running

| 3277008
| NFS Kerberos must be disabled on all LIFs of Vserver before deleting the
NFS configuration. When all LIFs are disabled, try the operation

| 3277009
| NFS Kerberos realms associated with the Vserver are deleted

| 3277111
| Internal error. Failed to remove NFS-specific security trace filter for
Vserver

```

```

| 3277112
| Internal error. Failed to modify the protocols field of a security trace
filter for Vserver
|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```



```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID3526cca40e353715d6f0c13181e32672]]
= Retrieve the NFS configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/services/{svm.uuid}`#
```

```
*Introduced In:* 9.6
```

Retrieves the NFS configuration of an SVM.

== Related ONTAP commands

```
* `vserver nfs show`
* `vserver nfs status`
```

== Learn more

```
* xref:{relative_path}protocols_nfs_services_endpoint_overview.html[DOC
/protocols/nfs/services]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|svm.uuid
```

```

|string
|path
|True
a|

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|auth_sys_extended_groups_enabled
|boolean
a|Specifies whether or not extended groups support over AUTH_SYS is
enabled.

|enabled
|boolean
a|Specifies if the NFS service is administratively enabled.

|extended_groups_limit
|integer
a|Specifies the maximum auxillary groups supported over AUTH_SYS and
RPCSEC_GSS.

|metric
|link:#metric[metric]

```

a|Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

|positive_cached_credential_ttl

|integer

a|Specifies the time to live value (in msec) of a positive cached credential

|protocol

|link:#protocol[protocol]

a|

|rquota_enabled

|boolean

a|Specifies whether or not the remote quota feature is enabled.

|showmount_enabled

|boolean

a|Specifies whether or not the showmount feature is enabled.

|state

|string

a|Specifies the state of the NFS service on the SVM. The following values are supported:

***** online - NFS server is ready to accept client requests.

***** offline - NFS server is not ready to accept client requests.

|statistics

|link:#statistics[statistics]

a|Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

|svm

|link:#svm[svm]

a|

|transport

|link:#transport[transport]

a|

```
|vstorage_enabled
|boolean
a|Specifies whether or not the VMware vstorage feature is enabled.
```

```
|===
```

```
.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "extended_groups_limit": "32",
  "metric": {
    "v3": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    }
  },
}
```

```

"v4": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"v41": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",

```

```

        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
}
},
"positive_cached_credential_ttl": "7200000",
"protocol": {
    "v4_id_domain": "string"
},
"state": "string",
"statistics": {
    "v3": {
        "iops_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "latency_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "status": "ok",
        "throughput_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "timestamp": "2017-01-25T11:20:13Z"
    },
    "v4": {
        "iops_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "latency_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "status": "ok",
        "throughput_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        }
    }
}

```

```

    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "v41": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  }
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|self  
|link:#href[href]  
a|
```

```
|===
```

```
[#iops]  
[.api-collapsible-fifth-title]  
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|other  
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read  
|integer  
a|Performance metric for read I/O operations.
```

```
|total  
|integer  
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#v3]
[.api-collapsible-fifth-title]
v3
```

The NFSv3 operations

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

|duration
|string
a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#v4]

[.api-collapsible-fifth-title]

v4

The NFSv4 operations

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|duration

|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops

|link:#iops[iops]

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#v41]
[.api-collapsible-fifth-title]
v41
```

The NFSv4.1 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|v3
|link:#v3[v3]
a|The NFSv3 operations
```



```
|v4
|link:#v4[v4]
a|The NFSv4 operations
```

```
|v41
|link:#v41[v41]
a|The NFSv4.1 operations
```

```
|===
```

```
[#v40_features]
[.api-collapsible-fifth-title]
v40_features
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|acl_enabled
```

```
|boolean
```

```
a|Specifies whether NFSv4.0 ACLs is enabled.
```

```
|read_delegation_enabled
```

```
|boolean
```

```
a|Specifies whether NFSv4.0 Read Delegation is enabled.
```

```
|write_delegation_enabled
```

```
|boolean
```

```
a|Specifies whether NFSv4.0 Write Delegation is enabled.
```

```
|===
```

```
[#v41_features]
[.api-collapsible-fifth-title]
v41_features
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|acl_enabled
|boolean
a|Specifies whether NFSv4.1 or later ACLs is enabled.

|pnfs_enabled
|boolean
a|Specifies whether NFSv4.1 or later Parallel NFS is enabled.

|read_delegation_enabled
|boolean
a|Specifies whether NFSv4.1 or later Read Delegation is enabled.

|write_delegation_enabled
|boolean
a|Specifies whether NFSv4.1 or later Write Delegation is enabled.

|===

[#protocol]
[.api-collapsible-fifth-title]
protocol

[cols=3*,options=header]
|===
|Name
|Type
|Description

|v3_64bit_identifiers_enabled
|boolean
a|Specifies whether 64-bit support for NFSv3 FSIDs and file IDs is
enabled.

|v3_enabled
|boolean
a|Specifies whether NFSv3 protocol is enabled.

```

```
|v40_enabled
|boolean
a|Specifies whether NFSv4.0 protocol is enabled.
```

```
|v40_features
|link:#v40_features[v40_features]
a|
```

```
|v41_enabled
|boolean
a|Specifies whether NFSv4.1 or later protocol is enabled.
```

```
|v41_features
|link:#v41_features[v41_features]
a|
```

```
|v4_64bit_identifiers_enabled
|boolean
a|Specifies whether 64-bit support for NFSv4.x FSIDs and file IDs is
enabled.
```

```
|v4_id_domain
|string
a|Specifies the domain portion of the string form of user and group
names as defined by the NFSv4 protocol.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
|===
|Name
```

```

|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#v3]
```

```
[.api-collapsible-fifth-title]
```

```
v3
```

The NFSv3 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#v4]
[.api-collapsible-fifth-title]
v4
```

The NFSv4 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.
```

```
|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
```

failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#v41]

[.api-collapsible-fifth-title]

v41

The NFSv4.1 operations

[cols=3*,options=header]

|===

|Name

|Type

|Description

|iops_raw

|link:#iops_raw[iops_raw]

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw

|link:#latency_raw[latency_raw]

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#statistics]

[.api-collapsible-fifth-title]

statistics

Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

[cols=3*,options=header]

|===

|Name

|Type

```

|Description

|v3
|link:#v3[v3]
a|The NFSv3 operations


|v4
|link:#v4[v4]
a|The NFSv4 operations


|v41
|link:#v41[v41]
a|The NFSv4.1 operations


|===


[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.


|uuid
|string
a|The unique identifier of the SVM.


|===

```

```

[#transport]
[.api-collapsible-fifth-title]
transport

[cols=3*,options=header]
|===
|Name
|Type
|Description

|tcp_enabled
|boolean
a|Specifies whether TCP transports are enabled on the server.

|udp_enabled
|boolean
a|Specifies whether UDP transports are enabled on the server.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]

```

```

[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID74334688ac830ff9281ccb5806c68886]]
= Update the NFS configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/nfs/services/{svm.uuid}`#

*Introduced In:* 9.6

Updates the NFS configuration of an SVM.

```

== Related ONTAP commands

```
* `vserver nfs modify`  
* `vserver nfs on`  
* `vserver nfs off`  
* `vserver nfs start`  
* `vserver nfs stop`
```

== Learn more

```
* xref:{relative_path}protocols_nfs_services_endpoint_overview.html[DOC  
/protocols/nfs/services]
```

== Parameters

[cols=5*,options=header]

|===

|Name
|Type
|In
|Required
|Description

|svm.uuid

|string

|path

|True

a|

|===

== Request Body

[cols=3*,options=header]

|===

|Name
|Type
|Description

|auth_sys_extended_groups_enabled

|boolean

a|Specifies whether or not extended groups support over AUTH_SYS is enabled.

```

|enabled
|boolean
a|Specifies if the NFS service is administratively enabled.

|extended_groups_limit
|integer
a|Specifies the maximum auxillary groups supported over AUTH_SYS and
RPCSEC_GSS.

|positive_cached_credential_ttl
|integer
a|Specifies the time to live value (in msecs) of a positive cached
credential

|protocol
|link:#protocol[protocol]
a|

|rquota_enabled
|boolean
a|Specifies whether or not the remote quota feature is enabled.

|showmount_enabled
|boolean
a|Specifies whether or not the showmount feature is enabled.

|state
|string
a|Specifies the state of the NFS service on the SVM. The following values
are supported:

***** online - NFS server is ready to accept client requests.

***** offline - NFS server is not ready to accept client requests.

|svm
|link:#svm[svm]
a|

|transport

```

```
|link:#transport[transport]
a|

|vstorage_enabled
|boolean
a|Specifies whether or not the VMware vstorage feature is enabled.

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "extended_groups_limit": "32",
  "positive_cached_credential_ttl": "7200000",
  "protocol": {
    "v4_id_domain": "string"
  },
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description
| 3276916
| Vserver is not running
```

```
| 3277069
| Cannot disable TCP because the SnapDiff RPC server is in the \"on\"
state

| 3277087
| Attempting to reduce the number of bits used for NFSv3 FSIDs and File
IDs from 64 to 32 on Vserver. This could result in collisions between
different File IDs and is not recommended

| 3277088
| Attempting to increase the number of bits used for NFSv3 FSIDs and File
IDs from 32 to 64 on Vserver. This could result in older client software
no longer working with the volumes owned by Vserver

| 3277090
| Attempting to disallow multiple FSIDs per mount point on Vserver. Since
this Vserver currently uses 32-bit NFSv3 FSIDs and File IDs, this could
result in collisions between different File IDs and is not recommended

| 3277099
| Domain name contains invalid characters or its too short. Allowed
characters are: alphabetical characters (A-Za-z), numeric characters (0-
9), minus sign (-), and the period (.). The first character must be
alphabetical or numeric, last character must not be a minus sign or a
period. Minimum supported length: 2 characters, maximum of 256 characters
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
```



```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

[.api-collapsible-fifth-title]

latency

The round trip latency in microseconds observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#throughput]

[.api-collapsible-fifth-title]

throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#v3]

[.api-collapsible-fifth-title]

v3

The NFSv3 operations

[cols=3*,options=header]

|===

|Name

|Type

|Description

|duration

|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops

|link:#iops[iops]

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#v4]

[.api-collapsible-fifth-title]

v4

The NFSv4 operations

[cols=3*,options=header]

|===

|Name

|Type

|Description

|duration

|string

a|The duration over which this sample is calculated. The time durations

are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#v41]
[.api-collapsible-fifth-title]
v41
```

The NFSv4.1 operations

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string

```

a|The timestamp of the performance data.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|v3

|link:#v3[v3]

a|The NFSv3 operations

|v4

|link:#v4[v4]

a|The NFSv4 operations

|v41

|link:#v41[v41]

a|The NFSv4.1 operations

|===

[#v40_features]
[.api-collapsible-fifth-title]
v40_features

[cols=3*,options=header]

|===

|Name
|Type
|Description

```
|acl_enabled
|boolean
a|Specifies whether NFSv4.0 ACLs is enabled.
```

```
|read_delegation_enabled
|boolean
a|Specifies whether NFSv4.0 Read Delegation is enabled.
```

```
|write_delegation_enabled
|boolean
a|Specifies whether NFSv4.0 Write Delegation is enabled.
```

```
|===
```

```
[#v41_features]
[.api-collapsible-fifth-title]
v41_features
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|acl_enabled
|boolean
a|Specifies whether NFSv4.1 or later ACLs is enabled.
```

```
|pnfs_enabled
|boolean
a|Specifies whether NFSv4.1 or later Parallel NFS is enabled.
```

```
|read_delegation_enabled
|boolean
a|Specifies whether NFSv4.1 or later Read Delegation is enabled.
```

```
|write_delegation_enabled
|boolean
a|Specifies whether NFSv4.1 or later Write Delegation is enabled.
```


|===

```
[#protocol]
[.api-collapsible-fifth-title]
protocol
```

[cols=3*,options=header]

|===

|Name

|Type

|Description

|v3_64bit_identifiers_enabled

|boolean

a|Specifies whether 64-bit support for NFSv3 FSIDs and file IDs is enabled.

|v3_enabled

|boolean

a|Specifies whether NFSv3 protocol is enabled.

|v40_enabled

|boolean

a|Specifies whether NFSv4.0 protocol is enabled.

|v40_features

|link:#v40_features[v40_features]

a|

|v41_enabled

|boolean

a|Specifies whether NFSv4.1 or later protocol is enabled.

|v41_features

|link:#v41_features[v41_features]

a|

|v4_64bit_identifiers_enabled

|boolean

a|Specifies whether 64-bit support for NFSv4.x FSIDs and file IDs is

enabled.

|v4_id_domain
|string
a|Specifies the domain portion of the string form of user and group names as defined by the NFSv4 protocol.

|===

[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|other
|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read
|integer

a|Performance metric for read I/O operations.

|total
|integer

a|Performance metric aggregated over all types of I/O operations.

|write
|integer

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]  
[.api-collapsible-fifth-title]  
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]  
[.api-collapsible-fifth-title]  
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#v3]
```

```
[.api-collapsible-fifth-title]
```

```
v3
```

The NFSv3 operations

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

```
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
```

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#v4]
[.api-collapsible-fifth-title]
v4

The NFSv4 operations

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.

|timestamp
|string
a|The timestamp of the performance data.

|===

```

```
[#v41]
[.api-collapsible-fifth-title]
v41
```

The NFSv4.1 operations

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.
```

```
|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
```

along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

Realtime performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|v3
|link:#v3[v3]
a|The NFSv3 operations
```

```
|v4
|link:#v4[v4]
a|The NFSv4 operations
```

```
|v41
|link:#v41[v41]
a|The NFSv4.1 operations
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```



```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#transport]
[.api-collapsible-fifth-title]
transport

[cols=3*,options=header]
|===
|Name
|Type
|Description

|tcp_enabled
|boolean
a|Specifies whether TCP transports are enabled on the server.

|udp_enabled
|boolean
a|Specifies whether UDP transports are enabled on the server.

|===

[#nfs_service]
[.api-collapsible-fifth-title]
nfs_service

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|auth_sys_extended_groups_enabled
|boolean
a|Specifies whether or not extended groups support over AUTH_SYS is
enabled.

|enabled
|boolean
a|Specifies if the NFS service is administratively enabled.

|extended_groups_limit
|integer
a|Specifies the maximum auxillary groups supported over AUTH_SYS and
RPCSEC_GSS.

|positive_cached_credential_ttl
|integer
a|Specifies the time to live value (in msecs) of a positive cached
credential

|protocol
|link:#protocol[protocol]
a|

|rquota_enabled
|boolean
a|Specifies whether or not the remote quota feature is enabled.

|showmount_enabled
|boolean
a|Specifies whether or not the showmount feature is enabled.

|state
|string
a|Specifies the state of the NFS service on the SVM. The following values
are supported:

```

```

***** online - NFS server is ready to accept client requests.

***** offline - NFS server is not ready to accept client requests.


|svm
|link:#svm[svm]
a|

|transport
|link:#transport[transport]
a|

|vstorage_enabled
|boolean
a|Specifies whether or not the VMware vstorage feature is enabled.


|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===

[#error]
[.api-collapsible-fifth-title]

```

```

error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====


[[IDef1bd0f7506ed92314b757451775c64f]]
= Retrieve NFS protocol historical performance metrics


[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nfs/services/{svm.uuid}/metrics`#


*Introduced In:* 9.7


Retrieves historical performance metrics for the NFS protocol of an SVM.

```

```

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|v4.latency.total
|integer
|query
|False
a|Filter by v4.latency.total

* Introduced in: 9.8

|v4.latency.read
|integer
|query
|False
a|Filter by v4.latency.read

* Introduced in: 9.8

|v4.latency.other
|integer
|query
|False
a|Filter by v4.latency.other

* Introduced in: 9.8

|v4.latency.write
|integer
|query
|False
a|Filter by v4.latency.write

* Introduced in: 9.8

```

```
|v4.status
|string
|query
|False
a|Filter by v4.status
```

* Introduced in: 9.8

```
|v4.iops.total
|integer
|query
|False
a|Filter by v4.iops.total
```

* Introduced in: 9.8

```
|v4.iops.read
|integer
|query
|False
a|Filter by v4.iops.read
```

* Introduced in: 9.8

```
|v4.iops.other
|integer
|query
|False
a|Filter by v4.iops.other
```

* Introduced in: 9.8

```
|v4.iops.write
|integer
|query
|False
a|Filter by v4.iops.write
```

* Introduced in: 9.8

```
|v4.duration
|string
```

```

|query
|False
a|Filter by v4.duration

* Introduced in: 9.8


|v4.throughput.write
|integer
|query
|False
a|Filter by v4.throughput.write

* Introduced in: 9.8


|v4.throughput.read
|integer
|query
|False
a|Filter by v4.throughput.read

* Introduced in: 9.8


|v4.throughput.total
|integer
|query
|False
a|Filter by v4.throughput.total

* Introduced in: 9.8


|v41.iops.total
|integer
|query
|False
a|Filter by v41.iops.total

* Introduced in: 9.8


|v41.iops.read
|integer
|query
|False

```

a|Filter by v41.iops.read

* Introduced in: 9.8

|v41.iops.other

|integer

|query

|False

a|Filter by v41.iops.other

* Introduced in: 9.8

|v41.iops.write

|integer

|query

|False

a|Filter by v41.iops.write

* Introduced in: 9.8

|v41.status

|string

|query

|False

a|Filter by v41.status

* Introduced in: 9.8

|v41.latency.total

|integer

|query

|False

a|Filter by v41.latency.total

* Introduced in: 9.8

|v41.latency.read

|integer

|query

|False

a|Filter by v41.latency.read

* Introduced in: 9.8

|v41.latency.other
|integer
|query
|False
a|Filter by v41.latency.other

* Introduced in: 9.8

|v41.latency.write
|integer
|query
|False
a|Filter by v41.latency.write

* Introduced in: 9.8

|v41.throughput.write
|integer
|query
|False
a|Filter by v41.throughput.write

* Introduced in: 9.8

|v41.throughput.read
|integer
|query
|False
a|Filter by v41.throughput.read

* Introduced in: 9.8

|v41.throughput.total
|integer
|query
|False
a|Filter by v41.throughput.total

* Introduced in: 9.8

```
|v41.duration
|string
|query
|False
a|Filter by v41.duration
```

* Introduced in: 9.8

```
|v3.duration
|string
|query
|False
a|Filter by v3.duration
```

```
|v3.throughput.write
|integer
|query
|False
a|Filter by v3.throughput.write
```

```
|v3.throughput.read
|integer
|query
|False
a|Filter by v3.throughput.read
```

```
|v3.throughput.total
|integer
|query
|False
a|Filter by v3.throughput.total
```

```
|v3.latency.total
|integer
|query
|False
a|Filter by v3.latency.total
```

```
|v3.latency.read
|integer
```

```
|query
|False
a|Filter by v3.latency.read

|v3.latency.other
|integer
|query
|False
a|Filter by v3.latency.other

|v3.latency.write
|integer
|query
|False
a|Filter by v3.latency.write

|v3.iops.total
|integer
|query
|False
a|Filter by v3.iops.total

|v3.iops.read
|integer
|query
|False
a|Filter by v3.iops.read

|v3.iops.other
|integer
|query
|False
a|Filter by v3.iops.other

|v3.iops.write
|integer
|query
|False
a|Filter by v3.iops.write
```

```
|v3.status
|string
|query
|False
a|Filter by v3.status
```

```
|timestamp
|string
|query
|False
a|Filter by timestamp
```

```
|svm.uuid
|string
|path
|True
a|Unique identifier of the SVM.
```

```
|interval
|string
|query
|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:
```

- * 1h: Metrics over the most recent hour sampled over 15 seconds.
- * 1d: Metrics over the most recent day sampled over 5 minutes.
- * 1w: Metrics over the most recent week sampled over 30 minutes.
- * 1m: Metrics over the most recent month sampled over 2 hours.
- * 1y: Metrics over the most recent year sampled over a day.
- * Default value: 1
- * enum: ["1h", "1d", "1w", "1m", "1y"]

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

```

* Max value: 120
* Min value: 0

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records
```

```
|records
|array[link:#records[records]]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "v3": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "latency": {
          "read": "200",
```

```

        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"v4": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "duration": "PT15S",
    "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"v41": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "duration": "PT15S",
    "iops": {
        "read": "200",
        "total": "1000",

```

```

        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    }
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {

```



```

        "code": "string",
        "message": "string"
    }
],
"code": "4",
"message": "entry doesn't exist",
"target": "uuid"
}
}
=====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next

```

```
|link:href[href]
```

```
a|
```

```
|self
```

```
|link:href[href]
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:href[href]
```

```
a|
```

```
|===
```

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#v3]
[.api-collapsible-fifth-title]
v3
```

The NFSv3 operations

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string

```

a|The timestamp of the performance data.

|===

[#v4]

[.api-collapsible-fifth-title]

v4

The NFSv4 operations

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|duration

|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops

|link:#iops[iops]

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with

"backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#v41]

[.api-collapsible-fifth-title]

v41

The NFSv4.1 operations

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|duration

|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops

|link:#iops[iops]

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#records]

[.api-collapsible-fifth-title]

records

Historical performance numbers, such as IOPS latency and throughput, for SVM-NFS protocol.

[cols=3*,options=header]

|===

|Name


```

|Type
|Description

|v3
|link:#v3[v3]
a|The NFSv3 operations

|v4
|link:#v4[v4]
a|The NFSv4 operations

|v41
|link:#v41[v41]
a|The NFSv4.1 operations

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= View and create Vscan configuration

:leveloffset: +1

[[IDecf9901f08a3eab27bf62945ddbc23ae]]
= Protocols Vscan endpoint overview

```

== Overview

Use Vscan to protect data from being compromised by viruses or other malicious code. Vscan combines best-in-class third party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when. Storage systems offload scanning operations to external servers hosting antivirus software from third party vendors. An Antivirus Connector on the external server handles communications between the storage system and the antivirus software.

== Examples

=== Retrieving all of the Vscan configurations

The API:

/api/protocols/vscan

The call:

```
curl -X GET "https://<mgmt-  
ip>/api/protocols/vscan?fields=*&return_records=true&return_timeout=15" -H  
"accept: application/json"
```

The response:

```
{  
  "records": [  
    {  
      "svm": {  
        "uuid": "03ce5c36-f269-11e8-8852-0050568e5298",  
        "name": "vs1"  
      },  
      "enabled": true,  
      "scanner_pools": [  
        {  
          "name": "scanner-1",  
          "servers": [  
            "1.1.1.1",  
            "10.72.204.27"  
          ],  
          "privileged_users": [  
            "cifs\\u1",  
            "cifs\\u2"  
          ],  
          "role": "primary",
```

```

    "cluster": {
      "name": "Cluster1",
      "uuid": "0228714d-f268-11e8-8851-0050568e5298"
    }
  },
  {
    "name": "scanner-2",
    "servers": [
      "1.1.1.1",
      "10.72.204.27"
    ],
    "privileged_users": [
      "cifs\\u1",
      "cifs\\u2"
    ],
    "role": "primary",
    "cluster": {
      "name": "Cluster1",
      "uuid": "0228714d-f268-11e8-8851-0050568e5298"
    }
  }
],
"on_access_policies": [
  {
    "name": "default_CIFS",
    "vsName": "vs1",
    "enabled": true,
    "mandatory": true,
    "scope": {
      "max_file_size": 2147483648,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    }
  },
  {
    "name": "on-access-test1",
    "vsName": "vs1",
    "enabled": false,
    "mandatory": true,
    "scope": {
      "max_file_size": 10000,
      "exclude_paths": [

```

```

        "\\dir"
    ],
    "include_extensions": [
        "mp*",
        "txt"
    ],
    "exclude_extensions": [
        "mp*",
        "txt"
    ],
    "scan_without_extension": true,
    "scan_readonly_volumes": false,
    "only_execute_access": false
}
},
{
    "name": "on-access-test2",
    "vsName": "vs1",
    "enabled": false,
    "mandatory": true,
    "scope": {
        "max_file_size": 10000,
        "exclude_paths": [
            "\\dir"
        ],
        "include_extensions": [
            "mp*",
            "txt"
        ],
        "exclude_extensions": [
            "mp*",
            "txt"
        ],
        "scan_without_extension": true,
        "scan_readonly_volumes": false,
        "only_execute_access": false
    }
}
],
"on_demand_policies": [
    {
        "name": "task-1",
        "scan_paths": [
            "/vol1"
        ],
        "log_path": "/vol1",

```

```

    "scope": {
      "max_file_size": 10000,
      "exclude_paths": [
        "/vol1"
      ],
      "include_extensions": [
        "vmdk",
        "mp*"
      ],
      "exclude_extensions": [
        "mp3",
        "mp4"
      ],
      "scan_without_extension": true
    }
  },
  {
    "name": "task-2",
    "scan_paths": [
      "/vol1"
    ],
    "log_path": "/vol2",
    "scope": {
      "max_file_size": 10000,
      "exclude_paths": [
        "/vol2"
      ],
      "include_extensions": [
        "vmdk",
        "mp*"
      ],
      "exclude_extensions": [
        "mp3",
        "mp4"
      ],
      "scan_without_extension": true
    }
  }
],
},
{
  "svm": {
    "uuid": "24c2567a-f269-11e8-8852-0050568e5298",
    "name": "vs2"
  },
  "enabled": false,

```

```

"scanner_pools": [
  {
    "name": "sp2",
    "servers": [
      "1.1.1.1"
    ],
    "privileged_users": [
      "cifs\\u1"
    ],
    "role": "idle"
  }
],
"on_access_policies": [
  {
    "name": "default_CIFS",
    "vsName": "vs2",
    "enabled": true,
    "mandatory": true,
    "scope": {
      "max_file_size": 2147483648,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    }
  },
  {
    "name": "ap1",
    "vsName": "vs2",
    "enabled": false,
    "mandatory": true,
    "scope": {
      "max_file_size": 2147483648,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    }
  }
],
"on_demand_policies": [
  {

```

```

        "name": "t1",
        "scan_paths": [
            "/vol1"
        ],
        "log_path": "/vol1",
        "scope": {
            "max_file_size": 10737418240,
            "include_extensions": [
                "*"
            ],
            "scan_without_extension": true
        }
    }
]
}
],
"num_records": 2
}
----

```

=== Retrieving all Vscan configurations for a particular SVM

The API:

/api/protocols/vscan/{svm.uuid}

The call:

```
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/24c2567a-f269-11e8-8852-0050568e5298?fields=*" -H "accept: application/json"
```

The response:

```

{
  "svm": {
    "uuid": "24c2567a-f269-11e8-8852-0050568e5298",
    "name": "vs2"
  },
  "enabled": false,
  "scanner_pools": [
    {
      "name": "sp2",
      "servers": [
        "1.1.1.1"
      ],
      "privileged_users": [
        "cifs\\u1"
      ]
    }
  ]
}

```



```

    ],
    "role": "idle"
  }
],
"on_access_policies": [
  {
    "name": "default_CIFS",
    "vsName": "vs2",
    "enabled": true,
    "mandatory": true,
    "scope": {
      "max_file_size": 2147483648,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    }
  },
  {
    "name": "ap1",
    "vsName": "vs2",
    "enabled": false,
    "mandatory": true,
    "scope": {
      "max_file_size": 2147483648,
      "include_extensions": [
        "*"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    }
  }
],
"on_demand_policies": [
  {
    "name": "t1",
    "scan_paths": [
      "/vol1"
    ],
    "log_path": "/vol1",
    "scope": {
      "max_file_size": 10737418240,
      "include_extensions": [

```

```

        "*"
    ],
    "scan_without_extension": true
}
}
]
}
----

=== Creating a Vscan configuration

----

# The API:
/api/protocols/vscan

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"enabled\": true, \"on_access_policies\": [ { \"enabled\": true,
  \"mandatory\": true, \"name\": \"on-access-test\", \"scope\": {
  \"exclude_extensions\": [ \"mp*\", \"txt\" ], \"exclude_paths\": [
  \"\\\\\\\\vol\" ], \"include_extensions\": [ \"mp*\", \"txt\" ],
  \"max_file_size\": 21474, \"only_execute_access\": false,
  \"scan_readonly_volumes\": false, \"scan_without_extension\": true } } ],
  \"on_demand_policies\": [ { \"log_path\": \"vol\", \"name\": \"task-1\",
  \"scan_paths\": [ \"vol\" ], \"schedule\": { \"name\": \"daily\",
  \"uuid\": \"d4984822-17b7-11e9-b450-0050568ecd85\" }, \"scope\": {
  \"exclude_extensions\": [ \"mp3\", \"mp4\" ], \"exclude_paths\": [
  \"vol\" ], \"include_extensions\": [ \"vmdk\", \"mp*\" ],
  \"max_file_size\": 10737, \"scan_without_extension\": true } } ],
  \"scanner_pools\": [ { \"cluster\": { \"name\": \"Cluster1\", \"uuid\":
  \"ab746d77-17b7-11e9-b450-0050568ecd85\" }, \"name\": \"scanner-1\",
  \"privileged_users\": [ \"cifs\\\\\\\\u1\", \"cifs\\\\\\\\u2\" ], \"role\":
  \"primary\", \"servers\": [ \"1.1.1.1\", \"10.72.204.27\" ] } ], \"svm\":
  { \"name\": \"vs1\", \"uuid\": \"b103be27-17b8-11e9-b451-0050568ecd85\"
  } } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "b103be27-17b8-11e9-b451-0050568ecd85",
        "name": "vs1"
      }
    }
  ]
}

```

```

},
"enabled": true,
"scanner_pools": [
  {
    "name": "scanner-1",
    "servers": [
      "1.1.1.1",
      "10.72.204.27"
    ],
    "privileged_users": [
      "cifs\\u1",
      "cifs\\u2"
    ],
    "role": "primary",
    "cluster": {
      "name": "Cluster1",
      "uuid": "ab746d77-17b7-11e9-b450-0050568ecd85"
    }
  }
],
"on_access_policies": [
  {
    "name": "on-access-test",
    "enabled": true,
    "mandatory": true,
    "scope": {
      "max_file_size": 21474,
      "exclude_paths": [
        "\\vol"
      ],
      "include_extensions": [
        "mp*",
        "txt"
      ],
      "exclude_extensions": [
        "mp*",
        "txt"
      ],
      "scan_without_extension": true,
      "scan_readonly_volumes": false,
      "only_execute_access": false
    }
  }
],
"on_demand_policies": [
  {

```

```

    "name": "task-1",
    "scan_paths": [
        "/vol"
    ],
    "log_path": "/vol",
    "schedule": {
        "uuid": "d4984822-17b7-11e9-b450-0050568ecd85",
        "name": "daily"
    },
    "scope": {
        "max_file_size": 10737,
        "exclude_paths": [
            "/"
        ],
        "include_extensions": [
            "vmdk",
            "mp*"
        ],
        "exclude_extensions": [
            "mp3",
            "mp4"
        ],
        "scan_without_extension": true
    }
}
]
}
]
}
----

=== Creating multiple Vscan scanner-pools for the specified SVM

----

# The API:
/api/protocols/vscan

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"scanner_pools\": [ { \"cluster\": { \"name\": \"Cluster1\", \"uuid\":
  \"ab746d77-17b7-11e9-b450-0050568ecd85\" }, \"name\": \"scanner-1\",
  \"privileged_users\": [ \"cifs\\\\\\\\u1\", \"cifs\\\\\\\\u2\" ], \"role\":
  \"primary\", \"servers\": [ \"1.1.1.1\", \"10.72.204.27\" ] }, {
  \"cluster\": { \"name\": \"Cluster1\", \"uuid\": \"ab746d77-17b7-11e9-

```

```
b450-0050568ecd85\" }, \"name\": \"scanner-2\", \"privileged_users\": [
  \"cifs\\\\\\u3\", \"cifs\\\\\\u4\" ], \"role\": \"primary\", \"servers\": [
  \"1.1.1.5\", \"10.72.3.27\" ] } ], \"svm\": { \"name\": \"vs1\", \"uuid\":
  \"b103be27-17b8-11e9-b451-0050568ecd85\" }}"
```

The response:

```
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "b103be27-17b8-11e9-b451-0050568ecd85",
        "name": "vs1"
      },
      "scanner_pools": [
        {
          "name": "scanner-1",
          "servers": [
            "1.1.1.1",
            "10.72.204.27"
          ],
          "privileged_users": [
            "cifs\\\\u1",
            "cifs\\\\u2"
          ],
          "role": "primary",
          "cluster": {
            "name": "Cluster1",
            "uuid": "ab746d77-17b7-11e9-b450-0050568ecd85"
          }
        },
        {
          "name": "scanner-2",
          "servers": [
            "1.1.1.5",
            "10.72.3.27"
          ],
          "privileged_users": [
            "cifs\\\\u3",
            "cifs\\\\u4"
          ],
          "role": "primary",
          "cluster": {
            "name": "Cluster1",
            "uuid": "ab746d77-17b7-11e9-b450-0050568ecd85"
          }
        }
      ]
    }
  ]
}
```

```

    }
  ]
}
]
}
----

=== Creating multiple Vscan On-access policies for a specified SVM

----

# The API:
/api/protocols/vscan

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"on_access_policies\": [ { \"enabled\": false, \"mandatory\": true,
  \"name\": \"on-access-test11\", \"scope\": { \"exclude_extensions\": [
  \"mp*\", \"txt\" ], \"exclude_paths\": [ \"\\\\\\\\vol\" ],
  \"include_extensions\": [ \"mp*\", \"txt\" ], \"max_file_size\": 214748,
  \"only_execute_access\": false, \"scan_readonly_volumes\": false,
  \"scan_without_extension\": true } }, { \"enabled\": false, \"mandatory\":
true, \"name\": \"on-access-test10\", \"scope\": { \"exclude_extensions\":
[ \"mp*\", \"txt\" ], \"exclude_paths\": [ \"\\\\\\\\vol\" ],
  \"include_extensions\": [ \"mp*\", \"txt\" ], \"max_file_size\": 21474,
  \"only_execute_access\": false, \"scan_readonly_volumes\": false,
  \"scan_without_extension\": true } } ], \"svm\": { \"name\": \"vs1\",
  \"uuid\": \"b103be27-17b8-11e9-b451-0050568ecd85\" } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "b103be27-17b8-11e9-b451-0050568ecd85",
        "name": "vs1"
      },
      "on_access_policies": [
        {
          "name": "on-access-test11",
          "enabled": false,
          "mandatory": true,
          "scope": {
            "max_file_size": 214748,

```

```

        "exclude_paths": [
            "\\vol"
        ],
        "include_extensions": [
            "mp*",
            "txt"
        ],
        "exclude_extensions": [
            "mp*",
            "txt"
        ],
        "scan_without_extension": true,
        "scan_readonly_volumes": false,
        "only_execute_access": false
    }
},
{
    "name": "on-access-test10",
    "enabled": false,
    "mandatory": true,
    "scope": {
        "max_file_size": 21474,
        "exclude_paths": [
            "\\vol"
        ],
        "include_extensions": [
            "mp*",
            "txt"
        ],
        "exclude_extensions": [
            "mp*",
            "txt"
        ],
        "scan_without_extension": true,
        "scan_readonly_volumes": false,
        "only_execute_access": false
    }
}
]
}
}
-----

```

=== Creating multiple Vscan On-demand policies for a specified SVM

```

----

# The API:
/api/protocols/vscan

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan?return_records=true"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"on_demand_policies\": [ { \"log_path\": \"/vol\", \"name\": \"task-1\",
  \"scan_paths\": [ \"/vol\" ], \"schedule\": { \"name\": \"daily\",
  \"uuid\": \"d4984822-17b7-11e9-b450-0050568ecd85\" }, \"scope\": {
  \"exclude_extensions\": [ \"mp3\", \"mp4\" ], \"exclude_paths\": [
  \"/vol1\" ], \"include_extensions\": [ \"vmdk\", \"mp*\" ],
  \"max_file_size\": 107374, \"scan_without_extension\": true } }, {
  \"log_path\": \"/vol\", \"name\": \"task-2\", \"scan_paths\": [ \"/vol\"
  ], \"scope\": { \"exclude_extensions\": [ \"mp3\", \"mp4\" ],
  \"exclude_paths\": [ \"/vol1\" ], \"include_extensions\": [ \"vmdk\",
  \"mp*\" ], \"max_file_size\": 107374, \"scan_without_extension\": true } }
  ], \"svm\": { \"name\": \"vs1\", \"uuid\": \"b103be27-17b8-11e9-b451-
  0050568ecd85\" } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "b103be27-17b8-11e9-b451-0050568ecd85",
        "name": "vs1"
      },
      "on_demand_policies": [
        {
          "name": "task-1",
          "scan_paths": [
            "/vol"
          ],
          "log_path": "/vol",
          "schedule": {
            "uuid": "d4984822-17b7-11e9-b450-0050568ecd85",
            "name": "daily"
          },
          "scope": {
            "max_file_size": 107374,
            "exclude_paths": [
              "/vol1"
            ],

```



```

        "include_extensions": [
            "vmdk",
            "mp*"
        ],
        "exclude_extensions": [
            "mp3",
            "mp4"
        ],
        "scan_without_extension": true
    }
},
{
    "name": "task-2",
    "scan_paths": [
        "/vol"
    ],
    "log_path": "/vol",
    "scope": {
        "max_file_size": 107374,
        "exclude_paths": [
            "/vol1"
        ],
        "include_extensions": [
            "vmdk",
            "mp*"
        ],
        "exclude_extensions": [
            "mp3",
            "mp4"
        ],
        "scan_without_extension": true
    }
}
]
}
}

```

=== Enabling Vscan for a specified SVM

```

# The API:
/api/protocols/vscan/{svm.uuid}

```

```
# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/03ce5c36-f269-11e8-8852-0050568e5298" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"enabled\": true}"
----
```

=== Clearing the Vscan cache for the specified SVM

```
-----

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/03ce5c36-f269-11e8-8852-0050568e5298" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"cache_clear\": true}"
-----
```

=== Deleting the Vscan configuration for a specified SVM

```
-----

# The API:
/api/protocols/vscan/{svm.uuid}
```

```
# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/vscan/03ce5c36-f269-11e8-8852-0050568e5298" -H "accept: application/json"
-----
```

[[ID754bababd7734b7dc403c55cebfcb253]]

= Retrieve the Vscan configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/vscan`#

Introduced In: 9.6

Retrieves the Vscan configuration.

This includes scanner-pools, On-Access policies, On-Demand policies, and information about whether a Vscan is enabled or disabled on an SVM.

Important notes:

* You can enable only one Vscan configuration at a time for an SVM.

* You can only query using `svm.uuid` or `svm.name`.

== Related ONTAP commands

- * `vserver vscan show`
- * `vserver vscan scanner-pool show`
- * `vserver vscan scanner-pool servers show`
- * `vserver vscan scanner-pool privileged-users show`
- * `vserver vscan scanner-pool show-active`
- * `vserver vscan on-access-policy show`
- * `vserver vscan on-access-policy file-ext-to-exclude show`
- * `vserver vscan on-access-policy file-ext-to-include show`
- * `vserver vscan on-access-policy paths-to-exclude show`
- * `vserver vscan on-demand-task show`

== Learn more

- * [xref:{relative_path}protocols_vscan_endpoint_overview.html\[DOC /protocols/vscan\]](#)
- * [xref:{relative_path}protocols_vscan_svm.uuid_scanner-pools_endpoint_overview.html\[DOC /protocols/vscan/{svm.uuid}/scanner-pools\]](#)

== Parameters

[cols=5*,options=header]
|===

Name	Type	In	Required	Description
------	------	----	----------	-------------

on_access_policies.name	string	query	False	a Filter by on_access_policies.name
-------------------------	--------	-------	-------	-------------------------------------

on_access_policies.scope.only_execute_access	boolean	query	False	a Filter by on_access_policies.scope.only_execute_access
----------------------------------------------	---------	-------	-------	----------------------------------------------------------

```
|on_access_policies.scope.include_extensions
|string
|query
|False
a|Filter by on_access_policies.scope.include_extensions

|on_access_policies.scope.exclude_extensions
|string
|query
|False
a|Filter by on_access_policies.scope.exclude_extensions

|on_access_policies.scope.scan_readonly_volumes
|boolean
|query
|False
a|Filter by on_access_policies.scope.scan_readonly_volumes

|on_access_policies.scope.exclude_paths
|string
|query
|False
a|Filter by on_access_policies.scope.exclude_paths

|on_access_policies.scope.max_file_size
|integer
|query
|False
a|Filter by on_access_policies.scope.max_file_size

|on_access_policies.scope.scan_without_extension
|boolean
|query
|False
a|Filter by on_access_policies.scope.scan_without_extension

|on_access_policies.mandatory
|boolean
|query
```

```

|False
a|Filter by on_access_policies.mandatory

|on_access_policies.enabled
|boolean
|query
|False
a|Filter by on_access_policies.enabled

|enabled
|boolean
|query
|False
a|Filter by enabled

|on_demand_policies.scan_paths
|string
|query
|False
a|Filter by on_demand_policies.scan_paths

|on_demand_policies.log_path
|string
|query
|False
a|Filter by on_demand_policies.log_path

|on_demand_policies.scope.include_extensions
|string
|query
|False
a|Filter by on_demand_policies.scope.include_extensions

|on_demand_policies.scope.scan_without_extension
|boolean
|query
|False
a|Filter by on_demand_policies.scope.scan_without_extension

|on_demand_policies.scope.max_file_size

```

```
|integer
|query
|False
a|Filter by on_demand_policies.scope.max_file_size

|on_demand_policies.scope.exclude_paths
|string
|query
|False
a|Filter by on_demand_policies.scope.exclude_paths

|on_demand_policies.scope.exclude_extensions
|string
|query
|False
a|Filter by on_demand_policies.scope.exclude_extensions

|on_demand_policies.schedule.name
|string
|query
|False
a|Filter by on_demand_policies.schedule.name

|on_demand_policies.schedule.uuid
|string
|query
|False
a|Filter by on_demand_policies.schedule.uuid

|on_demand_policies.name
|string
|query
|False
a|Filter by on_demand_policies.name

|scanner_pools.cluster.name
|string
|query
|False
a|Filter by scanner_pools.cluster.name
```

```
|scanner_pools.cluster.uuid  
|string  
|query  
|False  
a|Filter by scanner_pools.cluster.uuid
```

```
|scanner_pools.role  
|string  
|query  
|False  
a|Filter by scanner_pools.role
```

```
|scanner_pools.privileged_users  
|string  
|query  
|False  
a|Filter by scanner_pools.privileged_users
```

```
|scanner_pools.name  
|string  
|query  
|False  
a|Filter by scanner_pools.name
```

```
|scanner_pools.servers  
|string  
|query  
|False  
a|Filter by scanner_pools.servers
```

```
|svm.uuid  
|string  
|query  
|False  
a|Filter by svm.uuid
```

```
|svm.name  
|string  
|query  
|False
```

a|Filter by svm.name

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|return_records

|boolean

|query

|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.


```
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|num_records
```

```
|integer
```

```
a|Number of records
```

```
|records
```

```
|array[link:#vscan[vscan]]
```

```
a|
```

```
|===
```

.Example response

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "_links": {
```

```
    "next": {
```

```
      "href": "/api/resourcelink"
```

```
    },
```

```
    "self": {
```

```
      "href": "/api/resourcelink"
```

```
    }
```

```
  },
```

```
  "records": [
```

```
    {
```

```
      "_links": {
```

```
        "self": {
```

```

        "href": "/api/resourcelink"
    }
},
"on_access_policies": [
    {
        "name": "on-access-test",
        "scope": {
            "exclude_extensions": [
                "mp*",
                "txt"
            ],
            "exclude_paths": [
                "\\dir1\\dir2\\name",
                "\\vol\\a b",
                "\\vol\\a,b\\"
            ],
            "include_extensions": [
                "mp*",
                "txt"
            ],
            "max_file_size": "2147483648"
        }
    }
],
"on_demand_policies": [
    {
        "log_path": "/vol0/report_dir",
        "name": "task-1",
        "scan_paths": [
            "/vol1/",
            "/vol2/cifs/"
        ],
        "schedule": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "weekly",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "scope": {
            "exclude_extensions": [
                "mp3",
                "mp4"
            ],

```

```

        "exclude_paths": [
            "/vol1/cold-files/",
            "/vol1/cifs/names"
        ],
        "include_extensions": [
            "vmdk",
            "mp*"
        ],
        "max_file_size": "10737418240"
    }
},
"scanner_pools": [
    {
        "cluster": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "name": "scanner-1",
        "privileged_users": [
            "cifs\\u1",
            "cifs\\u2"
        ],
        "role": "string",
        "servers": [
            "1.1.1.1",
            "10.72.204.27",
            "vmwin204-27.fsct.nb"
        ]
    }
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}

```

```
]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|self
|link:#href[href]
a|

|===

[#scope]
[.api-collapsible-fifth-title]
scope

[cols=3*,options=header]
|===
|Name
|Type
|Description

|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions
|array[string]
a|List of file extensions to be scanned.

|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.

|only_execute_access
|boolean
a|Scan only files opened with execute-access.

|scan_readonly_volumes
|boolean
a|Specifies whether or not read-only volume can be scanned.

```

```
|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.
```

```
|===
```

```
[#vscan_on_access]
[.api-collapsible-fifth-title]
vscan_on_access
```

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|enabled
|boolean
```

```
a|Status of the On-Access Vscan policy
```

```
|mandatory
|boolean
```

```
a|Specifies if scanning is mandatory. File access is denied if there are
no external virus-scanning servers available for virus scanning.
```

```
|name
|string
a|On-Access policy ame
```

```
|scope
|link:#scope[scope]
a|
```

```
|===
```

```
[#schedule]
[.api-collapsible-fifth-title]
schedule
```

Schedule of the task.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|Job schedule name
```

```
|uuid
|string
a|Job schedule UUID
```

```
|===
```

```
[#scope]
[.api-collapsible-fifth-title]
scope
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.
```



```

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions
|array[string]
a|List of file extensions to be scanned.

|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.

|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.

|===

[#vscan_on_demand]
[.api-collapsible-fifth-title]
vscan_on_demand

Use On-Demand scanning to check files for viruses on a schedule. An On-
Demand policy defines the scope of an On-Demand scan.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|log_path
|string
a|The path from the Vserver root where the task report is created.

|name
|string
a|On-Demand task name

```

```
|scan_paths
|array[string]
a|List of paths that need to be scanned.
```

```
|schedule
|link:#schedule[schedule]
a|Schedule of the task.
```

```
|scope
|link:#scope[scope]
a|
```

```
|===
```

```
[#cluster_reference]
[.api-collapsible-fifth-title]
cluster_reference
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
 |_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#vscan_scanner_pool]
[.api-collapsible-fifth-title]
vscan_scanner_pool
```

Scanner pool is a set of attributes which are used to validate and manage

connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cluster
```

```
|link:#cluster_reference[cluster_reference]
```

```
a|
```

```
|name
```

```
|string
```

a|Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".".

```
|privileged_users
```

```
|array[string]
```

a|Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.

```
* example: ["cifs\u1", "cifs\u2"]
```

```
* Introduced in: 9.6
```

```
* x-omitempty: true
```

```
|role
```

```
|string
```

a|Specifies the role of the scanner pool. The possible values are:

```
*** primary    - Always active.
```

```
*** secondary - Active only when none of the primary external virus-scanning servers are connected.
```

```
*** idle       - Always inactive.
```

```
|servers
|array[string]
a|Specifies a list of IP addresses or FQDN for each Vscan server host
names which are allowed to connect to clustered ONTAP.
```

```
* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]
* Introduced in: 9.6
* x-omitempty: true
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

SVM, applies only to SVM-scoped objects.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#vscan]
[.api-collapsible-fifth-title]
vscan
```

Vscan can be used to protect data from being compromised by viruses or other malicious code. This combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when. Storage systems offload scanning operations to external servers hosting antivirus software from thirdparty vendors. An Antivirus Connector on the external server handles communications between the storage system and the antivirus software.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|cache_clear
```

```
|boolean
```

```
a|Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only
```

```
|enabled
```

```
|boolean
```

```
a|Specifies whether or not Vscan is enabled on the SVM.
```

```
|on_access_policies
```

```
|array[link:#vscan_on_access[vscan_on_access]]
```

```
a|
```

```
|on_demand_policies
```

```
|array[link:#vscan_on_demand[vscan_on_demand]]
```

```
a|
```

```
|scanner_pools
```

```
|array[link:#vscan_scanner_pool[vscan_scanner_pool]]
```

```
a|
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|SVM, applies only to SVM-scoped objects.
```

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

|===

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID39647c2a4421737e84d607bbe3dcd764]]
= Create a Vscan configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/vscan`#

*Introduced In:* 9.6

Creates a Vscan configuration, which includes a list of scanner-pools,
Vscan On-Access policies and Vscan On-Demand policies. Defines whether the
Vscan configuration you create is enabled or disabled for a specified SVM.

Important notes:

* You can enable only one Vscan configuration at a time for an SVM.
* There needs to be at least one active scanner-pool and one enabled On-
Access policy to enable Vscan successfully.
* By default, a Vscan is enabled when it's created.
* By default, the Vscan On-Access policies created from this endpoint are
in the disabled state. You can use the On-Access policy PATCH endpoint to
enable a particular On-Access policy. In ONTAP 9.6, only one Vscan On-
Access policy can be enabled and only one Vscan On-Demand policy can be
scheduled on an SVM.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the Vscan
configuration.

```

== Recommended optional properties

* ``scanner_pools`` - There must be at least one active scanner-pool for Vscan configuration. Created either through Vscan POST operation or scanner-pools POST operation.

== Default property values

If not specified in POST, the following default property value is assigned:

* ``enabled`` - `_true_`

== Related ONTAP commands

```
* `vserver vscan enable`
* `vserver vscan scanner-pool create`
* `vserver vscan scanner-pool apply-policy`
* `vserver vscan scanner-pool servers add`
* `vserver vscan scanner-pool privileged-users add`
* `vserver vscan on-access-policy create`
* `vserver vscan on-access-policy file-ext-to-exclude add`
* `vserver vscan on-access-policy file-ext-to-include add`
* `vserver vscan on-access-policy paths-to-exclude add`
* `vserver vscan on-demand-task create`
```

== Learn more

```
* xref:{relative_path}protocols_vscan_endpoint_overview.html[DOC
/protocols/vscan]
* xref:{relative_path}protocols_vscan_svm.uuid_scanner-
pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-
pools]
```

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records


```

|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cache_clear
|boolean
a|Discards the cached information of the files that have been successfully
scanned. Once the cache is cleared, files are scanned again when they are
accessed. PATCH only

|enabled
|boolean
a|Specifies whether or not Vscan is enabled on the SVM.

|on_access_policies
|array[link:#vscan_on_access[vscan_on_access]]
a|

|on_demand_policies
|array[link:#vscan_on_demand[vscan_on_demand]]
a|

|scanner_pools
|array[link:#vscan_scanner_pool[vscan_scanner_pool]]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.

```

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "on_access_policies": [
    {
      "name": "on-access-test",
      "scope": {
        "exclude_extensions": [
          "mp*",
          "txt"
        ],
        "exclude_paths": [
          "\\dir1\\dir2\\name",
          "\\vol\\a b",
          "\\vol\\a,b\\"
        ],
        "include_extensions": [
          "mp*",
          "txt"
        ],
        "max_file_size": "2147483648"
      }
    }
  ],
  "on_demand_policies": [
    {
      "log_path": "/vol0/report_dir",
      "name": "task-1",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scope": {
        "exclude_extensions": [
          "mp3",
          "mp4"
        ]
      }
    }
  ]
}
```

```

    ],
    "exclude_paths": [
        "/vol1/cold-files/",
        "/vol1/cifs/names"
    ],
    "include_extensions": [
        "vmdk",
        "mp*"
    ],
    "max_file_size": "10737418240"
}
},
"scanner_pools": [
{
    "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "name": "scanner-1",
    "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
    ],
    "role": "string",
    "servers": [
        "1.1.1.1",
        "10.72.204.27",
        "vmwin204-27.fsct.nb"
    ]
}
},
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Response

```

Status: 201, Created

```

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|num_records
|integer
a|Number of records

|records
|array[link:#vscan[vscan]]
a|

|===

```

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```

{
  "records": [
    {
      "on_access_policies": [
        {
          "name": "on-access-test",
          "scope": {
            "exclude_extensions": [
              "mp*",
              "txt"
            ],
            "exclude_paths": [
              "\\dir1\\dir2\\name",
              "\\vol\\a b",
              "\\vol\\a,b\\"
            ],
            "include_extensions": [
              "mp*",
              "txt"
            ],
            "max_file_size": "2147483648"
          }
        }
      ],
      "on_demand_policies": [
        {

```

```

    "log_path": "/vol0/report_dir",
    "name": "task-1",
    "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
    ],
    "schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "scope": {
        "exclude_extensions": [
            "mp3",
            "mp4"
        ],
        "exclude_paths": [
            "/vol1/cold-files/",
            "/vol1/cifs/names"
        ],
        "include_extensions": [
            "vmdk",
            "mp*"
        ],
        "max_file_size": "10737418240"
    }
},
"scanner_pools": [
    {
        "cluster": {
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "name": "scanner-1",
        "privileged_users": [
            "cifs\\u1",
            "cifs\\u2"
        ],
        "role": "string",
        "servers": [
            "1.1.1.1",
            "10.72.204.27",
            "vmwin204-27.fsct.nb"
        ]
    }
],

```

```

    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
]
}
====

```

== Error

Status: Default

ONTAP Error Response Codes

//start table

[cols=2*,options=header]

|===

//header

| Error Code | Description

//end header

//end row

//start row

|10027259 +

//end row

//start row

|A scanner-pool, an On-Access policy, or an On-Demand policy might fail to get created due to either a systematic error or some hardware failure. The error code returned details the failure along with the reason for the failure. For example, if a scanner-pool fails due to an incorrect cluster name, then the error might read: "Failed to create scanner-pool "scanner-1". Reason: "Cluster uuid points to different cluster name instead of the cluster-name supplied.". Retry the operation."

//end row

//start row

|10027260 +

//end row

//start row

|If a scanner-pool, an On-Access policy or an On-Demand policy specified in the input already exists, then a duplicate error is returned. For example, if a scanner-pool "scanner-1" already exists for an SVM and is again specified in the input, the error message will read: " Failed to create scanner-pool "scanner-1" as the specified entry already exists. Delete the entry and retry the POST operation."

//end row

//start row

```

|2621462 +
//end row
//start row
|The specified SVM name is invalid
//end row
//start row
|2621706 +
//end row
//start row
|The specified svm.uuid is either invalid or belongs to a different SVM
//end row
//start row
|10027015 +
//end row
//start row
|Attempting to enable a Vscan but no active scanner-pool exists for the
specified SVM
//end row
//start row
|10027011 +
//end row
//start row
|Attempting to enable a Vscan for an SVM for which no CIFS server exists
//end row
//start row
|10027023 +
//end row
//start row
|Attempting to enable a Vscan for an SVM for which no active Vscan On-
Access policy exist
//end row
|===
//end table

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#scope]
[.api-collapsible-fifth-title]
scope

[cols=3*,options=header]
|===
|Name
|Type
|Description

|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions
|array[string]
a|List of file extensions to be scanned.

|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.

|only_execute_access

```



```
|boolean
a|Scan only files opened with execute-access.
```

```
|scan_readonly_volumes
|boolean
a|Specifies whether or not read-only volume can be scanned.
```

```
|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.
```

```
|===
```

```
[#vscan_on_access]
[.api-collapsible-fifth-title]
vscan_on_access
```

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|enabled
|boolean
a|Status of the On-Access Vscan policy
```

```
|mandatory
|boolean
a|Specifies if scanning is mandatory. File access is denied if there are no external virus-scanning servers available for virus scanning.
```

```
|name
|string
a|On-Access policy ame
```

```

|scope
|link:#scope[scope]
a|

|===

[#schedule]
[.api-collapsible-fifth-title]
schedule

Schedule of the task.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name


|uuid
|string
a|Job schedule UUID


|===


[#scope]
[.api-collapsible-fifth-title]
scope


[cols=3*,options=header]
|===
|Name
|Type
|Description

|exclude_extensions
|array[string]

```

a|List of file extensions for which scanning is not performed.

|exclude_paths

|array[string]

a|List of file paths for which scanning must not be performed.

|include_extensions

|array[string]

a|List of file extensions to be scanned.

|max_file_size

|integer

a|Maximum file size, in bytes, allowed for scanning.

|scan_without_extension

|boolean

a|Specifies whether or not files without any extension can be scanned.

|===

[#vscan_on_demand]

[.api-collapsible-fifth-title]

vscan_on_demand

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|log_path

|string

a|The path from the Vserver root where the task report is created.

|name

|string

a|On-Demand task name

|scan_paths

|array[string]

a|List of paths that need to be scanned.

|schedule

|link:#schedule[schedule]

a|Schedule of the task.

|scope

|link:#scope[scope]

a|

|===

[#cluster_reference]

[.api-collapsible-fifth-title]

cluster_reference

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|

|uuid

|string

a|

|===

[#vscan_scanner_pool]

[.api-collapsible-fifth-title]

vscan_scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or

```
"Vscan server".
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cluster
```

```
|link:#cluster_reference[cluster_reference]
```

```
a|
```

```
|name
```

```
|string
```

a|Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".".

```
|privileged_users
```

```
|array[string]
```

a|Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remediating and quarantining operations.

```
* example: ["cifs\u1", "cifs\u2"]
```

```
* Introduced in: 9.6
```

```
* x-omitempty: true
```

```
|role
```

```
|string
```

a|Specifies the role of the scanner pool. The possible values are:

```
*** primary    - Always active.
```

```
*** secondary - Active only when none of the primary external virus-scanning servers are connected.
```

```
*** idle       - Always inactive.
```

```
|servers
```

```
|array[string]
a|Specifies a list of IP addresses or FQDN for each Vscan server host
names which are allowed to connect to clustered ONTAP.

* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]
* Introduced in: 9.6
* x-omitempty: true
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

SVM, applies only to SVM-scoped objects.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#vscan]
[.api-collapsible-fifth-title]
vscan
```

Vscan can be used to protect data from being compromised by viruses or other malicious code. This combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when. Storage systems offload scanning operations to external servers hosting antivirus software from thirdparty

vendors. An Antivirus Connector on the external server handles communications between the storage system and the antivirus software.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cache_clear
```

```
|boolean
```

```
a|Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only
```

```
|enabled
```

```
|boolean
```

```
a|Specifies whether or not Vscan is enabled on the SVM.
```

```
|on_access_policies
```

```
|array[link:#vscan_on_access[vscan_on_access]]
```

```
a|
```

```
|on_demand_policies
```

```
|array[link:#vscan_on_demand[vscan_on_demand]]
```

```
a|
```

```
|scanner_pools
```

```
|array[link:#vscan_scanner_pool[vscan_scanner_pool]]
```

```
a|
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|SVM, applies only to SVM-scoped objects.
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

```
error_arguments
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```



```
|===
```

```
//end collapsible .Definitions block  
=====
```

```
:leveloffset: -1
```

```
= Manage Vscan configuration
```

```
:leveloffset: +1
```

```
[[ID0bc97120663389e1dcf2a60cecdf27f0]]  
= Protocols Vscan server-status endpoint overview
```

```
== Overview
```

This API is used to display connection status information for the external virus-scanning servers or "Vscan servers".

```
== Examples
```

```
=== Retrieving all fields for the Vscan server status
```

```
'''
```

```
----
```

```
# The API:  
/api/protocols/vscan/server_status/
```

```
# The call:  
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/server_status?fields=*"   
-H "accept: application/hal+json"
```

```
# The response:
```

```
{  
  "records": [  
    {  
      "svm": {
```

```

    "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
      }
    }
  },
  "node": {
    "uuid": "fe696362-f138-11e8-8e22-0050568e0945",
    "name": "Cluster-01",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/fe696362-f138-11e8-8e22-
0050568e0945"
      }
    }
  },
  "ip": "10.141.46.173",
  "type": "primary",
  "state": "disconnected",
  "disconnected_reason": "unknown",
  "_links": {
    "self": {
      "href": "/api/protocols/vscan/server_status/86fbc414-f140-11e8-
8e22-0050568e0945/Cluster-01/10.141.46.173"
    }
  }
},
{
  "svm": {
    "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
      }
    }
  },
  "node": {
    "uuid": "fe696362-f138-11e8-8e22-0050568e0945",
    "name": "Cluster-01",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/fe696362-f138-11e8-8e22-
0050568e0945"
      }
    }
  }
}

```

```

    }
  }
},
"ip": "fd20:8b1e:b255:5053::46:173",
"type": "primary",
"state": "disconnected",
"disconnected_reason": "remote_closed",
"_links": {
  "self": {
    "href": "/api/protocols/vscan/server_status/86fbc414-f140-11e8-8e22-0050568e0945/Cluster-01/fd20%3A8b1e%3Ab255%3A5053%3A%3A46%3A173"
  }
}
},
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/vscan/server_status?fields=*"
  }
}
}
}

```

'''

=== Retrieving the server status information for the server with IP address 10.141.46.173

'''

The API:
/api/protocols/vscan/server_status

The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/server_status?ip=10.141.46.173&fields=*" -H "accept: application/hal+json"

The response:
{
"records": [
 {
 "svm": {

```

    "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
      }
    }
  },
  "node": {
    "uuid": "fe696362-f138-11e8-8e22-0050568e0945",
    "name": "Cluster-01",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/fe696362-f138-11e8-8e22-
0050568e0945"
      }
    }
  },
  "ip": "10.141.46.173",
  "type": "primary",
  "state": "connected",
  "update_time": "2018-12-19T08:03:40.988Z",
  "vendor": "XYZ",
  "version": "1.12.2",
  "_links": {
    "self": {
      "href": "/api/protocols/vscan/server_status/86fbc414-f140-11e8-
8e22-0050568e0945/Cluster-01/10.141.46.173"
    }
  }
},
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/protocols/vscan/server_status?ip=10.141.46.173&fields=*"
  }
}
}
----

'''

```

```
[[ID30740b445d8bd763c0da6cd141c91d63]]
```

= Retrieve the Vscan server status

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/vscan/server-status`#
```

Introduced In: 9.6

Retrieves a Vscan server status.

== Related ONTAP commands

* `vserver vscan connection-status show-all`

== Learn more

* xref:{relative_path}protocols_vscan_server-status_endpoint_overview.html[DOC /protocols/vscan/server-status]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|node.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by node.name
```

```
|node.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by node.uuid
```

```
|svm.uuid
```

```
|string
```

```
|query  
|False  
a|Filter by svm.uuid
```

```
|svm.name  
|string  
|query  
|False  
a|Filter by svm.name
```

```
|vendor  
|string  
|query  
|False  
a|Filter by vendor
```

```
|state  
|string  
|query  
|False  
a|Filter by state
```

```
|ip  
|string  
|query  
|False  
a|Filter by ip
```

```
|disconnected_reason  
|string  
|query  
|False  
a|Filter by disconnected_reason
```

```
|version  
|string  
|query  
|False  
a|Filter by version
```

```
|update_time
|string
|query
|False
a|Filter by update_time
```

```
|type
|string
|query
|False
a|Filter by type
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#vscan_server_status[vscan_server_status]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {

```



```

    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "disconnected_reason": "string",
      "ip": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "update_time": "string",
      "vendor": "string",
      "version": "string"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]

```

node

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

SVM, applies only to SVM-scoped objects.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

a|The unique identifier of the SVM.

|===

```
[#vscan_server_status]
[.api-collapsible-fifth-title]
vscan_server_status
```

Displays the connection status information of the external virus-scanning servers.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|disconnected_reason
|string
```

a|Specifies the server disconnected reason.

The following is a list of the possible reasons:

* unknown	- Disconnected, unknown reason.
* vscan_disabled	- Disconnected, Vscan is disabled on the SVM.
* no_data_lif	- Disconnected, SVM does not have data LIF.
* session_uninitialized	- Disconnected, session is not initialized.
* remote_closed	- Disconnected, server has closed the connection.
* invalid_protocol_msg_received.	- Disconnected, invalid protocol message received.
* invalid_session_id	- Disconnected, invalid session ID received.
* inactive_connection	- Disconnected, no activity on connection.
* invalid_user	- Connection request by an invalid user.
* server_removed	- Disconnected, server has been removed from the active Scanners List.

enum:

- * unknown
- * vscan_disabled
- * no_data_lif
- * session_uninitialized
- * remote_closed
- * invalid_protocol_msg
- * invalid_session_id

```
* inactive_connection
```

```
* invalid_user
```

```
* server_removed
```

```
|ip
```

```
|string
```

```
a|IP address of the Vscan server.
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|state
```

```
|string
```

```
a|Specifies the server connection state indicating if it is in the  
connected or disconnected state.
```

```
The following is a list of the possible states:
```

```
* connected                - Connected
```

```
* disconnected              - Disconnected
```

```
enum:
```

```
* connected
```

```
* disconnected
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|SVM, applies only to SVM-scoped objects.
```

```
|type
```

```
|string
```

```
a|Server type. The possible values are:
```

```
*** primary - Primary server
```

```
*** backup  - Backup server
```

```
|update_time
```

```
|string
```

```
a|Specifies the time the server is in the connected or disconnected state.
```

```
|vendor
```

```
|string  
a|Name of the connected virus-scanner vendor.
```

```
|version  
|string  
a|Version of the connected virus-scanner.
```

```
|===
```

```
[#error_arguments]  
[.api-collapsible-fifth-title]  
error_arguments
```

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|arguments  
|array[link:#error_arguments[error_arguments]]
```

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[ID7158762738823a18d134868a965ea755]]

= Delete a Vscan configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/protocols/vscan/{svm.uuid}`#

Introduced In: 9.6

Deletes a Vscan configuration.

Important notes:

* The Vscan DELETE endpoint deletes all of the Vscan configuration of an SVM. It first disables the Vscan and then deletes all of the SVM scanner-pools, On-Access policies, and On-Demand policies.

* Disable the active Vscan On-Access policy on an SVM before performing the Vscan delete operation on that SVM.

== Related ONTAP commands

* `vserver vscan scanner-pool delete`


```
* `vserver vscan on-access-policy delete`
* `vserver vscan on-demand-policy delete`

== Learn more

* xref:{relative_path}protocols_vscan_endpoint_overview.html[DOC
/protocols/vscan]
* xref:{relative_path}protocols_vscan_svm.uuid_scanner-
pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-
pools]
```

== Parameters

```
[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===
```

== Response

Status: 200, Ok

== Error

Status: Default

```
ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error Code | Description
```

```
//end header
//end row
//start row
|10027259 +
//end row
//start row
|A scanner-pool, an On-Access policy, or an On-Demand policy might fail to
get deleted due to either a systematic error or some hardware failure. The
error code returned details the failure along with the reason for the
failure. For example, "Failed to delete On-Access policy "spl". Reason:
"Failed to delete policy. Reason: policy must be disabled before being
deleted.". Retry the operation."
//end row
|===
//end table
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
```

```

    }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[ID9744525addde4a343866bfd13ce761d9]]
= Retrieve the Vscan configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/vscan/{svm.uuid}`#
```

Introduced In: 9.6

Retrieves the Vscan configuration for a specified SVM.
This includes scanner-pools, On-Access policies, On-Demand policies, and
information about whether a Vscan is enabled or disabled on an SVM.

Important note:

* You can enable only one Vscan configuration at a time for an SVM.

== Related ONTAP commands

```
* `vserver vscan show`
* `vserver vscan scanner-pool show`
* `vserver vscan scanner-pool servers show`
* `vserver vscan scanner-pool privileged-users show`
```

```
* `vserver vscan scanner-pool show-active`
* `vserver vscan on-access-policy show`
* `vserver vscan on-access-policy file-ext-to-exclude show`
* `vserver vscan on-access-policy file-ext-to-include show`
* `vserver vscan on-access-policy paths-to-exclude show`
* `vserver vscan on-demand-task show`
```

== Learn more

```
* xref:{relative_path}protocols_vscan_endpoint_overview.html[DOC
/protocols/vscan]
* xref:{relative_path}protocols_vscan_svm.uuid_scanner-
pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-
pools]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|===
```

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|cache_clear
|boolean
a|Discards the cached information of the files that have been successfully
scanned. Once the cache is cleared, files are scanned again when they are
accessed. PATCH only

|enabled
|boolean
a|Specifies whether or not Vscan is enabled on the SVM.

|on_access_policies
|array[link:#vscan_on_access[vscan_on_access]]
a|

|on_demand_policies
|array[link:#vscan_on_demand[vscan_on_demand]]
a|

|scanner_pools
|array[link:#vscan_scanner_pool[vscan_scanner_pool]]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{

```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "on_access_policies": [
      {
        "name": "on-access-test",
        "scope": {
          "exclude_extensions": [
            "mp*",
            "txt"
          ],
          "exclude_paths": [
            "\\dir1\\dir2\\name",
            "\\vol\\a b",
            "\\vol\\a,b\\"
          ],
          "include_extensions": [
            "mp*",
            "txt"
          ],
          "max_file_size": "2147483648"
        }
      }
    ],
    "on_demand_policies": [
      {
        "log_path": "/vol0/report_dir",
        "name": "task-1",
        "scan_paths": [
          "/vol1/",
          "/vol2/cifs/"
        ],
        "schedule": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "weekly",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "scope": {
          "exclude_extensions": [
            "mp3",

```

```

        "mp4"
    ],
    "exclude_paths": [
        "/vol1/cold-files/",
        "/vol1/cifs/names"
    ],
    "include_extensions": [
        "vmdk",
        "mp*"
    ],
    "max_file_size": "10737418240"
}
}
],
"scanner_pools": [
{
    "cluster": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "name": "scanner-1",
    "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
    ],
    "role": "string",
    "servers": [
        "1.1.1.1",
        "10.72.204.27",
        "vmwin204-27.fsct.nb"
    ]
}
],
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}

```



```

    }
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#scope]
[.api-collapsible-fifth-title]
scope

[cols=3*,options=header]
|===
|Name
|Type
|Description

|exclude_extensions
|array[string]

```

a|List of file extensions for which scanning is not performed.

|exclude_paths

|array[string]

a|List of file paths for which scanning must not be performed.

|include_extensions

|array[string]

a|List of file extensions to be scanned.

|max_file_size

|integer

a|Maximum file size, in bytes, allowed for scanning.

|only_execute_access

|boolean

a|Scan only files opened with execute-access.

|scan_readonly_volumes

|boolean

a|Specifies whether or not read-only volume can be scanned.

|scan_without_extension

|boolean

a|Specifies whether or not files without any extension can be scanned.

|===

[#vscan_on_access]

[.api-collapsible-fifth-title]

vscan_on_access

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

|enabled
|boolean
a|Status of the On-Access Vscan policy


|mandatory
|boolean
a|Specifies if scanning is mandatory. File access is denied if there are
no external virus-scanning servers available for virus scanning.


|name
|string
a|On-Access policy ame


|scope
|link:#scope[scope]
a|

|===


[#schedule]
[.api-collapsible-fifth-title]
schedule

Schedule of the task.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string

```

a|Job schedule name

|uuid

|string

a|Job schedule UUID

|===

[#scope]

[.api-collapsible-fifth-title]

scope

[cols=3*,options=header]

|===

|Name

|Type

|Description

|exclude_extensions

|array[string]

a|List of file extensions for which scanning is not performed.

|exclude_paths

|array[string]

a|List of file paths for which scanning must not be performed.

|include_extensions

|array[string]

a|List of file extensions to be scanned.

|max_file_size

|integer

a|Maximum file size, in bytes, allowed for scanning.

|scan_without_extension

|boolean

a|Specifies whether or not files without any extension can be scanned.

|===

```
[#vscan_on_demand]
[.api-collapsible-fifth-title]
vscan_on_demand
```

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|log_path
|string
a|The path from the Vserver root where the task report is created.

|name
|string
a|On-Demand task name

|scan_paths
|array[string]
a|List of paths that need to be scanned.

|schedule
|link:#schedule[schedule]
a|Schedule of the task.

|scope
|link:#scope[scope]
a|

|===

[#cluster_reference]
[.api-collapsible-fifth-title]
cluster_reference
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#vscan_scanner_pool]
```

```
[.api-collapsible-fifth-title]
```

```
vscan_scanner_pool
```

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cluster
```

```
|link:#cluster_reference[cluster_reference]
```

```
a|
```

```
|name
```

```
|string
```

a|Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".".

```

|privileged_users
|array[string]
a|Specifies a list of privileged users. A valid form of privileged user-
name is "domain-name\user-name". Privileged user-names are stored and
treated as case-insensitive strings. Virus scanners must use one of the
registered privileged users for connecting to clustered Data ONTAP for
exchanging virus-scanning protocol messages and to access file for
scanning, remedying and quarantining operations.

* example: ["cifs\u1", "cifs\u2"]
* Introduced in: 9.6
* x-omitempty: true


|role
|string
a|Specifies the role of the scanner pool. The possible values are:

*** primary    - Always active.

*** secondary - Active only when none of the primary external virus-
scanning servers are connected.

*** idle       - Always inactive.


|servers
|array[string]
a|Specifies a list of IP addresses or FQDN for each Vscan server host
names which are allowed to connect to clustered ONTAP.

* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]
* Introduced in: 9.6
* x-omitempty: true


|===


[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.

[cols=3*,options=header]

```



```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]

```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[IDd8e8b72cafe47339d3d66305ed6828db]]
```

```
= Update the Vscan configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-  
block]#`/protocols/vscan/{svm.uuid}`#
```

```
*Introduced In:* 9.6
```

Updates the Vscan configuration of an SVM. Allows you to either enable or disable a Vscan, and allows you to clear the Vscan cache that stores the past scanning data for an SVM.

Important note:

* The Vscan PATCH endpoint does not allow you to modify scanner-pools, On-Demand policies or On-Access policies. Those modifications can only be done through their respective endpoints.

== Related ONTAP commands

- * ``vserver vscan enable``
- * ``vserver vscan disable``
- * ``vserver vscan reset``

== Learn more

- * `xref:{relative_path}protocols_vscan_endpoint_overview.html[DOC /protocols/vscan]`
- * `xref:{relative_path}protocols_vscan_svm.uuid_scanner-pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-pools]`

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

svm.uuid
string
path
True
a UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===

```

|Name
|Type
|Description

|cache_clear
|boolean
a|Discards the cached information of the files that have been successfully
scanned. Once the cache is cleared, files are scanned again when they are
accessed. PATCH only

|enabled
|boolean
a|Specifies whether or not Vscan is enabled on the SVM.

|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error Code | Description
//end header
//end row
//start row
|10027015 +
//end row
//start row
|Attempting to enable a Vscan but no active scanner-pool exists for the
specified SVM
//end row
//start row
|10027011 +
//end row

```

```
//start row
|Attempting to enable a Vscan for an SVM for which no CIFS server exists
//end row
//start row
|10027023 +
//end row
//start row
|Attempting to enable a Vscan for an SVM for which no active Vscan On-
Access policy exists
//end row
|===
//end table
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

====

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|href
|string
a|
```

|===

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#scope]
[.api-collapsible-fifth-title]
scope
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions
|array[string]
a|List of file extensions to be scanned.

|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.

|only_execute_access
|boolean
a|Scan only files opened with execute-access.

|scan_readonly_volumes
|boolean
a|Specifies whether or not read-only volume can be scanned.

|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.

|===

[#vscan_on_access]
[.api-collapsible-fifth-title]
vscan_on_access

```

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|Status of the On-Access Vscan policy
```

```
|mandatory
```

```
|boolean
```

```
a|Specifies if scanning is mandatory. File access is denied if there are  
no external virus-scanning servers available for virus scanning.
```

```
|scope
```

```
|link:#scope[scope]
```

```
a|
```

```
|===
```

```
[#schedule]
```

```
[.api-collapsible-fifth-title]
```

```
schedule
```

```
Schedule of the task.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|Job schedule name
```

```
|uuid
|string
a|Job schedule UUID
```

```
|===
```

```
[#scope]
[.api-collapsible-fifth-title]
scope
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.
```

```
|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.
```

```
|include_extensions
|array[string]
a|List of file extensions to be scanned.
```

```
|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.
```

```
|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.
```

```
|===
```



```
[#vscan_on_demand]
[.api-collapsible-fifth-title]
vscan_on_demand
```

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|log_path
|string
a|The path from the Vserver root where the task report is created.
```

```
|scan_paths
|array[string]
a|List of paths that need to be scanned.
```

```
|schedule
|link:#schedule[schedule]
a|Schedule of the task.
```

```
|scope
|link:#scope[scope]
a|
```

```
|===
```

```
[#cluster_reference]
[.api-collapsible-fifth-title]
cluster_reference
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#vscan_scanner_pool]
```

```
[.api-collapsible-fifth-title]
```

```
vscan_scanner_pool
```

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cluster
```

```
|link:#cluster_reference[cluster_reference]
```

```
a|
```

```
|privileged_users
```

```
|array[string]
```

a|Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.

* example: ["cifs\u1", "cifs\u2"]

* Introduced in: 9.6

* x-omitempty: true

```
|role
```

```
|string
```

a|Specifies the role of the scanner pool. The possible values are:

*** primary - Always active.

*** secondary - Active only when none of the primary external virus-scanning servers are connected.

*** idle - Always inactive.

|servers

|array[string]

a|Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.

* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]

* Introduced in: 9.6

* x-omitempty: true

|===

[#svm]

[.api-collapsible-fifth-title]

svm

SVM, applies only to SVM-scoped objects.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

```
[#vscan]
[.api-collapsible-fifth-title]
vscan
```

Vscan can be used to protect data from being compromised by viruses or other malicious code. This combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when. Storage systems offload scanning operations to external servers hosting antivirus software from thirdparty vendors. An Antivirus Connector on the external server handles communications between the storage system and the antivirus software.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|cache_clear
|boolean
```

a|Discards the cached information of the files that have been successfully scanned. Once the cache is cleared, files are scanned again when they are accessed. PATCH only

```
|enabled
|boolean
```

a|Specifies whether or not Vscan is enabled on the SVM.

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
```

a|Argument code

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
:leveloffset: -1
```

```
= Manage Vscan On-Access policies
```

```
:leveloffset: +1
```

```
[[ID5583ec97186b6b2a0034f8551c518f54]]
```

```
= Protocols Vscan svm.uuid on-access-policies endpoint overview
```

```
== Overview
```

Use Vscan On-Access scanning to actively scan file objects for viruses when clients access files over SMB. To control which file operations trigger a vscan, use Vscan File-Operations Profile (vscan-fileop-profile) option in the CIFS share. The Vscan On-Access policy configuration defines the scope and status of On-Access scanning on file objects. Use this API to retrieve and manage Vscan On-Access policy configurations and Vscan On-Access policy statuses for the SVM.

```
== Examples
```

```
=== Retrieving all fields for all policies of an SVM
```

```
'''
```

```
----
```

```
# The API:
```

```
/api/protocols/vscan/{svm.uuid}/on-access-policies/
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/{svm.uuid}/on-access-policies?fields=*" -H "accept: application/hal+json"
```

```
# The response:
```

```
{
```

```
"records": [
```

```
{
```

```
  "svm": {
```

```
    "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
```

```
    "name": "vs1",
```

```
    "_links": {
```

```

        "self": {
            "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
    },
    "name": "default_CIFS",
    "enabled": true,
    "mandatory": true,
    "scope": {
        "max_file_size": 2147483648,
        "include_extensions": [
            "*"
        ],
        "scan_without_extension": true,
        "scan_readonly_volumes": false,
        "only_execute_access": false
    },
    "_links": {
        "self": {
            "href": "/api/protocols/vscan/179d3c85-7053-11e8-b9b8-005056b41bd1/on-access-policies/default_CIFS"
        }
    }
},
{
    "svm": {
        "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
        "name": "vs1",
        "_links": {
            "self": {
                "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
            }
        }
    },
    "name": "on-access-policy",
    "enabled": false,
    "mandatory": true,
    "scope": {
        "max_file_size": 3221225472,
        "exclude_paths": [
            "\\vol\\a b\\",
            "\\vol\\a,b\\"
        ],
        "include_extensions": [
            "mp*",
            "tx*"
        ]
    }
}

```

```

    ],
    "exclude_extensions": [
        "mp3",
        "txt"
    ],
    "scan_without_extension": true,
    "scan_readonly_volumes": false,
    "only_execute_access": true
},
"_links": {
    "self": {
        "href": "/api/protocols/vscan/179d3c85-7053-11e8-b9b8-005056b41bd1/on-access-policies/on-access-policy"
    }
}
],
"num_records": 2,
"_links": {
    "self": {
        "href": "/api/protocols/vscan/179d3c85-7053-11e8-b9b8-005056b41bd1/on-access-policies?fields=*"
    }
}
}
}
}
}
}

'''

=== Retrieving the specific On-Access policy associated with the specified SVM

'''

----

# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/179d3c85-7053-11e8-b9b8-005056b41bd1/on-access-policies/on-access-policy" -H "accept: application/json"

# The response:
{

```



```

"svm": {
  "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
  "name": "vs1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
},
"name": "on-access-policy",
"enabled": true,
"mandatory": true,
"scope": {
  "max_file_size": 3221225472,
  "exclude_paths": [
    "\\vol\\a b\\",
    "\\vol\\a,b\\"
  ],
  "include_extensions": [
    "mp*",
    "tx*"
  ],
  "exclude_extensions": [
    "mp3",
    "txt"
  ],
  "scan_without_extension": true,
  "scan_readonly_volumes": false,
  "only_execute_access": true
},
"_links": {
  "self": {
    "href": "/api/protocols/vscan/179d3c85-7053-11e8-b9b8-005056b41bd1/on-access-policies/task1"
  }
}
}
}

```

'''

=== Creating a Vscan On-Access policy

The Vscan On-Access policy POST endpoint creates an On-Access policy for the specified SVM. Set enabled to "true" to enable scanning on the created policy.

The API:

/api/protocols/vscan/{svm.uuid}/on-access-policies

The call:

```
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-access-policies?return_records=true" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"enabled\": false, \"mandatory\": true, \"name\": \"on-access-policy\", \"scope\": { \"exclude_extensions\": [ \"txt\", \"mp3\" ], \"exclude_paths\": [ \"\\\\\\\\dir1\\\\\\\\dir2\\\\\\\\ame\", \"\\\\\\\\vol\\\\\\\\a b\" ], \"include_extensions\": [ \"mp*\", \"txt\" ], \"max_file_size\": 3221225472, \"only_execute_access\": true, \"scan_readonly_volumes\": false, \"scan_without_extension\": true }}"
```

The response:

```
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "vs1"
      },
      "name": "on-access-policy",
      "enabled": false,
      "mandatory": true,
      "scope": {
        "max_file_size": 3221225472,
        "exclude_paths": [
          "\\dir1\\dir2\\ame",
          "\\vol\\a b"
        ],
        "include_extensions": [
          "mp*",
          "txt"
        ],
        "exclude_extensions": [
          "txt",
          "mp3"
        ],
        "scan_without_extension": true,
        "scan_readonly_volumes": false,
        "only_execute_access": true
      }
    }
  ]
}
```

```

    }
]
}
----

'''

=== Creating a Vscan On-Access policy where a number of optional fields
are not specified

'''

----

# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-access-policies?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"enabled\":
false, \"mandatory\": true, \"name\": \"on-access-policy\", \"scope\": {
\"exclude_paths\": [ \"\\\\\\\\vol\\\\\\\\a b\", \"\\\\\\\\vol\\\\\\\\a,b\\\\\\\\\" ],
\"max_file_size\": 1073741824, \"scan_without_extension\": true }}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "vs1"
      },
      "name": "on-access-policy",
      "enabled": false,
      "mandatory": true,
      "scope": {
        "max_file_size": 1073741824,
        "exclude_paths": [
          "\\vol\\a b",
          "\\vol\\a,b\\"
        ],
        "scan_without_extension": true
      }
    }
  ]
}

```

```

}
----

'''

=== Updating a Vscan On-Access policy

The policy being modified is identified by the UUID of the SVM and the
policy name.

----

# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies/{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-access-policies/on-access-policy" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"scope\":
{ \"include_extensions\": [ \"txt\" ], \"only_execute_access\": true,
\"scan_readonly_volumes\": false, \"scan_without_extension\": true }}"
----

'''

=== Deleting a Vscan On-Access policy

The policy to be deleted is identified by the UUID of the SVM and the
policy name.

----

# The API:
/api/protocols/vscan/{svm.uuid}/on-access-policies/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-
8e22-0050568e0945/on-access-policies/on-access-policy" -H "accept:
application/hal+json"
----

'''

```

```
[[ID88ae3e897afd12b7bf2c0ec9709879ce]]
```

= Retrieve a Vscan On-Access policy

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/vscan/{svm.uuid}/on-access-policies`#
```

Introduced In: 9.6

Retrieves the Vscan On-Access policy.

== Related ONTAP commands

* `vserver vscan on-access-policy show`

* `vserver vscan on-access-policy file-ext-to-include show`

* `vserver vscan on-access-policy file-ext-to-exclude show`

* `vserver vscan on-access-policy paths-to-exclude show`

== Learn more

* xref:{relative_path}protocols_vscan_svm.uuid_on-access-policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-access-policies]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by name
```

```
|scope.only_execute_access
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Filter by scope.only_execute_access
```

```
|scope.include_extensions
|string
|query
|False
a|Filter by scope.include_extensions
```

```
|scope.exclude_extensions
|string
|query
|False
a|Filter by scope.exclude_extensions
```

```
|scope.scan_readonly_volumes
|boolean
|query
|False
a|Filter by scope.scan_readonly_volumes
```

```
|scope.exclude_paths
|string
|query
|False
a|Filter by scope.exclude_paths
```

```
|scope.max_file_size
|integer
|query
|False
a|Filter by scope.max_file_size
```

```
|scope.scan_without_extension
|boolean
|query
|False
a|Filter by scope.scan_without_extension
```

```
|mandatory
|boolean
|query
```

```

|False
a|Filter by mandatory

|enabled
|boolean
|query
|False
a|Filter by enabled

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.

```

When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

- * Default value: 1
- * Max value: 120
- * Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records
```

```
|records
|array[link:#vscan_on_access[vscan_on_access]]
a|
```

```
|===
```

.Example response
[%collapsible%closed]


```

=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "name": "on-access-test",
      "scope": {
        "exclude_extensions": [
          "mp*",
          "txt"
        ],
        "exclude_paths": [
          "\\dir1\\dir2\\name",
          "\\vol\\a b",
          "\\vol\\a,b\\"
        ],
        "include_extensions": [
          "mp*",
          "txt"
        ],
        "max_file_size": "2147483648"
      }
    }
  ]
}
=====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#scope]
[.api-collapsible-fifth-title]
scope

[cols=3*,options=header]
|===
|Name
|Type
|Description

|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions
|array[string]
a|List of file extensions to be scanned.

```

```
|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.
```

```
|only_execute_access
|boolean
a|Scan only files opened with execute-access.
```

```
|scan_readonly_volumes
|boolean
a|Specifies whether or not read-only volume can be scanned.
```

```
|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.
```

```
|===
```

```
[#vscan_on_access]
[.api-collapsible-fifth-title]
vscan_on_access
```

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|enabled
|boolean
a|Status of the On-Access Vscan policy
```

```
|mandatory
```

```

|boolean
a|Specifies if scanning is mandatory. File access is denied if there are
no external virus-scanning servers available for virus scanning.

|name
|string
a|On-Access policy ame

|scope
|link:#scope[scope]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID42b8644aaf505b5b8648c9c52b7fe851]]
= Create a Vscan On-Access policy

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/vscan/{svm.uuid}/on-access-policies`#

*Introduced In:* 9.6

Creates a Vscan On-Access policy. Created only on a data SVM.
</b>Important notes:

* You must enable the policy on an SVM before its files can be scanned.
* You can enable only one On-Access policy at a time on an SVM. By
default, the policy is enabled on creation. * If the Vscan On-Access
policy has been created successfully on an SVM but cannot be enabled due
to an error, the Vscan On-Access policy configurations are saved. The

```

Vscan On-Access policy is then enabled using the PATCH operation.

== Required properties

- * ``svm.uuid`` - Existing SVM in which to create the Vscan On-Access policy.
- * ``name`` - Name of the Vscan On-Access policy. Maximum length is 256 characters.

== Default property values

If not specified in POST, the following default property values are assigned:

- * ``enabled`` - `_true_`
- * ``mandatory`` - `_true_`
- * ``include_extensions`` - `_*_`
- * ``max_file_size`` - `_2147483648_`
- * ``only_execute_access`` - `_false_`
- * ``scan_readonly_volumes`` - `_false_`
- * ``scan_without_extension`` - `_true_`

== Related ONTAP commands

- * ``vserver vscan on-access-policy create``
- * ``vserver vscan on-access-policy enable``
- * ``vserver vscan on-access-policy disable``
- * ``vserver vscan on-access-policy file-ext-to-include add``
- * ``vserver vscan on-access-policy file-ext-to-exclude add``
- * ``vserver vscan on-access-policy paths-to-exclude add``

== Learn more

* `xref:{relative_path}protocols_vscan_svm.uuid_on-access-policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-access-policies]`

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

```

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Status of the On-Access Vscan policy

|mandatory
|boolean
a|Specifies if scanning is mandatory. File access is denied if there are
no external virus-scanning servers available for virus scanning.

|name
|string
a|On-Access policy ame

|scope
|link:#scope[scope]
a|

```


|===

.Example request

[%collapsible%closed]

====

```
[source,json,subs=+macros]
{
  "name": "on-access-test",
  "scope": {
    "exclude_extensions": [
      "mp*",
      "txt"
    ],
    "exclude_paths": [
      "\\dir1\\dir2\\name",
      "\\vol\\a b",
      "\\vol\\a,b\\"
    ],
    "include_extensions": [
      "mp*",
      "txt"
    ],
    "max_file_size": "2147483648"
  }
}
```

====

== Response

Status: 201, Created

[cols=3*,options=header]

|===

|Name

|Type

|Description

|num_records

|integer

a|Number of records

|records

|array[link:#vscan_on_access[vscan_on_access]]

```

a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "name": "on-access-test",
      "scope": {
        "exclude_extensions": [
          "mp*",
          "txt"
        ],
        "exclude_paths": [
          "\\dir1\\dir2\\name",
          "\\vol\\a b",
          "\\vol\\a,b\\"
        ],
        "include_extensions": [
          "mp*",
          "txt"
        ],
        "max_file_size": "2147483648"
      }
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 10027043
| The new On-Access policy cannot be created as the SVM has reached the
maximum number of On-Access policies allowed. Delete an existing policy in

```

order to create a new policy

| 10027101

| The file size must be in the range 1KB to 1TB

| 10027107

| The include extensions list cannot be empty. Specify at least one extension for inclusion

| 10027109

| The specified CIFS path is invalid. It must be in the form "\dir1\dir2" or "\dir1\dir2\"

| 10027249

| The On-Access policy created successfully but failed to enable the policy. The reason for enable policy operation failure might be that another policy is enabled. Disable the enabled policy and then enable the newly created policy using the PATCH operation

| 10027253

| The number of paths specified exceeds the configured number of maximum paths. You cannot specify more than the maximum number of configured paths

| 10027254

| The number of extensions specified exceeds the configured maximum number of extensions. You cannot specify more than the maximum number of configured extensions

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#scope]

[.api-collapsible-fifth-title]

scope

[cols=3*,options=header]

|===

|Name

|Type

|Description

```

|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions
|array[string]
a|List of file extensions to be scanned.

|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.

|only_execute_access
|boolean
a|Scan only files opened with execute-access.

|scan_readonly_volumes
|boolean
a|Specifies whether or not read-only volume can be scanned.

|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.

|===

[#vscan_on_access]
[.api-collapsible-fifth-title]
vscan_on_access

```

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Status of the On-Access Vscan policy

|mandatory
|boolean
a|Specifies if scanning is mandatory. File access is denied if there are
no external virus-scanning servers available for virus scanning.

|name
|string
a|On-Access policy ame

|scope
|link:#scope[scope]
a|

|===

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message

```

```
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID02b0bea8bd24de2645a9ce722b76085d]]
= Delete an antivirus On-Access policy configuration
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/vscan/{svm.uuid}/on-access-policies/{name}`#
```

Introduced In: 9.6

Deletes the anti-virus On-Access policy configuration.

== Related ONTAP commands

* `vserver vscan on-access-policy delete`

== Learn more

* xref:{relative_path}protocols_vscan_svm.uuid_on-access-
policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-access-
policies]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
```

```

|Description

|name
|string
|path
|True
a|

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 10027034
| An On-Access policy associated with an administrative SVM cannot be
deleted.

| 10027040
| An On-Access policy with a status enabled cannot be deleted. Disable the
policy and then delete the policy.
|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string

```

```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

[[ID8b0e325d39c71987c4ac07420b14e54e]]

= Retrieve the Vscan On-Access policy configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/vscan/{svm.uuid}/on-access-policies/{name}`#

Introduced In: 9.6

Retrieves the Vscan On-Access policy configuration of an SVM.

== Related ONTAP commands

* `vserver vscan on-access-policy show`
* `vserver vscan on-access-policy file-ext-to-include show`
* `vserver vscan on-access-policy file-ext-to-exclude show`
* `vserver vscan on-access-policy paths-to-exclude show`

== Learn more

* xref:{relative_path}protocols_vscan_svm.uuid_on-access-policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-access-policies]

== Parameters

[cols=5*,options=header]

|===

|Name
|Type
|In
|Required
|Description

|name
|string
|path
|True
a|

|svm.uuid
|string
|path
|True

```
a|UUID of the SVM to which this object belongs.
```

```
|fields  
|array[string]  
|query  
|False  
a|Specify the fields to return.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|enabled  
|boolean
```

```
a|Status of the On-Access Vscan policy
```

```
|mandatory  
|boolean
```

```
a|Specifies if scanning is mandatory. File access is denied if there are  
no external virus-scanning servers available for virus scanning.
```

```
|name  
|string  
a|On-Access policy ame
```

```
|scope  
|link:#scope[scope]  
a|
```

```
|===
```

```
.Example response  
[%collapsible%closed]
```

```

=====
[source,json,subs=+macros]
{
  "name": "on-access-test",
  "scope": {
    "exclude_extensions": [
      "mp*",
      "txt"
    ],
    "exclude_paths": [
      "\\dir1\\dir2\\name",
      "\\vol\\a b",
      "\\vol\\a,b\\"
    ],
    "include_extensions": [
      "mp*",
      "txt"
    ],
    "max_file_size": "2147483648"
  }
}
=====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{

```

```

"error": {
  "arguments": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#scope]
[.api-collapsible-fifth-title]
scope

```

```

[cols=3*,options=header]

```

|===

```

|Name
|Type
|Description

```

```

|exclude_extensions

```

```

|array[string]

```

a|List of file extensions for which scanning is not performed.

```

|exclude_paths

```

```

|array[string]

```

a|List of file paths for which scanning must not be performed.

```

|include_extensions

```

```

|array[string]

```

a|List of file extensions to be scanned.

```

|max_file_size

```

```

|integer
a|Maximum file size, in bytes, allowed for scanning.

|only_execute_access
|boolean
a|Scan only files opened with execute-access.

|scan_readonly_volumes
|boolean
a|Specifies whether or not read-only volume can be scanned.

|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]

```

```

error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID3246e41a0fb4bcecea6eefalc7a5588f]]
= Update the Vscan On-Access policy configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/vscan/{svm.uuid}/on-access-policies/{name}`#

*Introduced In:* 9.6

```

Updates the Vscan On-Access policy configuration and/or enables/disables the Vscan On-Access policy of an SVM. You cannot modify the configurations for an On-Access policy associated with an administrative SVM, although

you can enable and disable the policy associated with an administrative SVM.

== Related ONTAP commands

```
* `vserver vscan on-access-policy modify`  
* `vserver vscan on-access-policy enable`  
* `vserver vscan on-access-policy disable`  
* `vserver vscan on-access-policy file-ext-to-include add`  
* `vserver vscan on-access-policy file-ext-to-exclude add`  
* `vserver vscan on-access-policy paths-to-exclude add`  
* `vserver vscan on-access-policy file-ext-to-include remove`  
* `vserver vscan on-access-policy file-ext-to-exclude remove`  
* `vserver vscan on-access-policy paths-to-exclude remove`
```

== Learn more

```
* xref:{relative_path}protocols_vscan_svm.uuid_on-access-  
policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-access-  
policies]
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|name  
|string  
|path  
|True  
a|
```

```
|svm.uuid  
|string  
|path  
|True  
a|UUID of the SVM to which this object belongs.
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|Status of the On-Access Vscan policy
```

```
|mandatory
```

```
|boolean
```

```
a|Specifies if scanning is mandatory. File access is denied if there are  
no external virus-scanning servers available for virus scanning.
```

```
|scope
```

```
|link:#scope[scope]
```

```
a|
```

```
|===
```

.Example request

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "scope": {
```

```
    "exclude_extensions": [
```

```
      "mp*",
```

```
      "txt"
```

```
    ],
```

```
    "exclude_paths": [
```

```
      "\\dir1\\dir2\\name",
```

```
      "\\vol\\a b",
```

```
      "\\vol\\a,b\\"
```

```
    ],
```

```
    "include_extensions": [
```

```
      "mp*",
```

```
      "txt"
```

```
    ],
```

```
"max_file_size": "2147483648"
}
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description

| 10027033
| Configurations for an On-Access policy associated with an administrative
SVM cannot be modified. However, the policy can be enabled or disabled.

| 10027046
| The specified SVM is not the owner of the specified policy. Check for
the correct SVM who owns the policy.

| 10027101
| The file size must be in the range 1KB to 1TB

| 10027107
| The include extensions list cannot be empty. Specify at least one
extension for inclusion.

| 10027109
| The specified CIFS path is invalid. It must be in the form "\dir1\dir2"
or "\dir1\dir2\".

| 10027249
| The On-Access policy updated successfully but failed to enable/disable
the policy. The reason for an enable policy operation failure might be
that another policy is enabled. Disable the already enabled policy and
then enable the policy. The reason for a disable policy operation failure
might be that Vscan is enabled on the SVM. Disable the Vscan first and
then disable the policy.
```

```

| 10027250
| The On-Access policy cannot be enabled/disabled. The reason for an
enable policy operation failure might be that another policy is enabled.
Disable the already enabled policy and then enable the policy. The reason
for a disable policy operation failure might be that Vscan is enabled on
the SVM. Disable the Vscan and then disable the policy.

| 10027253
| The number of paths specified exceeds the configured maximum number of
paths. You cannot specify more than the maximum number of configured
paths.

| 10027254
| The number of extensions specified exceeds the configured maximum number
of extensions. You cannot specify more than the maximum number of
configured extensions.
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#scope]
[.api-collapsible-fifth-title]
scope

[cols=3*,options=header]
|===
|Name
|Type
|Description

|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions

```

```

|array[string]
a|List of file extensions to be scanned.

|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.

|only_execute_access
|boolean
a|Scan only files opened with execute-access.

|scan_readonly_volumes
|boolean
a|Specifies whether or not read-only volume can be scanned.

|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.

```

```

|===

```

```

[#vscan_on_access]
[.api-collapsible-fifth-title]
vscan_on_access

```

An On-Access policy that defines the scope of an On-Access scan. Use On-Access scanning to check for viruses when clients open, read, rename, or close files over CIFS. By default, ONTAP creates an On-Access policy named "default_CIFS" and enables it for all the SVMs in a cluster.

```

[cols=3*,options=header]

```

```

|===

```

```

|Name
|Type
|Description

```

```

|enabled
|boolean
a|Status of the On-Access Vscan policy

```

```
|mandatory
|boolean
a|Specifies if scanning is mandatory. File access is denied if there are
no external virus-scanning servers available for virus scanning.
```

```
|scope
|link:#scope[scope]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
= Manage Vscan On-Demand policies
```

```
:leveloffset: +1
```

```
[[IDd5f00d8ef2b767cb993fe4dfe41460b3]]
= Protocols Vscan svm.uuid on-demand-policies endpoint overview
```

```
== Overview
```

Vscan On-Demand scanning is used to check files for viruses on a schedule. For example, it can be used to run scans only in off-peak hours, or to scan very large files that are excluded from an on-access scan. Vscan On-Demand scanning can be used for any path in the SVM namespace.

Vscan On-Demand policy configurations define the scope of a Vscan On-Demand scan. The schedule parameter in the On-Demand policy configuration decides when to execute the task. Schedule can be created using the /api/clusters/schedule endpoint and can be assigned on policy create or policy modify. This API is used to retrieve and manage Vscan On-Demand policy configurations. It is also used to schedule the Vscan On-Demand scan.

== Examples

=== Retrieving all fields for all policies of an SVM

'''

The API:

/api/protocols/vscan/{svm.uuid}/on-demand-policies/

The call:

curl -X GET "https://<mgmt-ip>/api/protocols/vscan/{svm.uuid}/on-demand-policies?fields=*" -H "accept: application/hal+json"

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
          }
        }
      },
      "name": "on-demand-policy1",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "log_path": "/vol0/report_dir",
      "schedule": {
        "uuid": "f6d0843e-f159-11e8-8e22-0050568e0945",
        "name": "schedule",
```



```

    "_links": {
      "self": {
        "href": "/api/cluster/schedules/f6d0843e-f159-11e8-8e22-0050568e0945"
      }
    },
    "scope": {
      "max_file_size": 10737418240,
      "exclude_paths": [
        "/vol1/cold-files/",
        "/vol1/cifs/names"
      ],
      "include_extensions": [
        "vmdk",
        "mp*"
      ],
      "exclude_extensions": [
        "mp3",
        "mp4"
      ],
      "scan_without_extension": false
    },
    "_links": {
      "self": {
        "href": "/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-demand-policies/policy1"
      }
    }
  },
  {
    "svm": {
      "uuid": "86fbc414-f140-11e8-8e22-0050568e0945",
      "name": "vs1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
        }
      }
    },
    "name": "on-demand-policy2",
    "scan_paths": [
      "/vol1/",
      "/vol2/cifs/"
    ],
    "log_path": "/report",

```



```

    "_links": {
      "self": {
        "href": "/api/svm/svms/86fbc414-f140-11e8-8e22-0050568e0945"
      }
    },
    "name": "on-demand-policy",
    "scan_paths": [
      "/vol1/cifs"
    ],
    "log_path": "/report",
    "scope": {
      "max_file_size": 10737418240,
      "include_extensions": [
        "vmdk",
        "mp*"
      ],
      "scan_without_extension": true
    },
    "_links": {
      "self": {
        "href": "/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-demand-policies/policy2"
      }
    }
  }
}

```

'''

=== Creating a Vscan On-Demand policy

The Vscan On-Demand policy POST endpoint creates an On-Demand policy for the specified SVM. Specify the schedule parameter to schedule an On-Demand scan.

The API:

/api/protocols/vscan/{svm.uuid}/on-demand-policies

The call:

```

curl -X POST "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-demand-policies?return_records=true" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"log_path\": \"/vol10/report_dir\", \"name\": \"on-demand-policy\", \"scan_paths\": [

```

```
\"/vol1/\", \"/vol2/cifs/\" ], \"schedule\": { \"name\": \"weekly\",
\"uuid\": \"1cd8a442-86d1-11e0-ae1c-123478563412\" }, \"scope\": {
\"exclude_extensions\": [ \"mp3\" ], \"exclude_paths\": [ \"/vol/cold-
files/\" ], \"include_extensions\": [ \"vmdk\", \"mp*\" ],
\"max_file_size\": 1073741824, \"scan_without_extension\": true }}"
```

The response:

```
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "vs1"
      },
      "name": "on-demand-policy",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "log_path": "/vol0/report_dir",
      "schedule": {
        "name": "weekly"
      },
      "scope": {
        "max_file_size": 1073741824,
        "exclude_paths": [
          "/vol/cold-files/"
        ],
        "include_extensions": [
          "vmdk",
          "mp*"
        ],
        "exclude_extensions": [
          "mp3"
        ],
        "scan_without_extension": true
      }
    }
  ]
}
----

'''
```

=== Creating a Vscan On-Demand policy where a number of optional fields are not specified

```

'''

----

# The API:
/api/protocols/vscan/{svm.uuid}/on-demand-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-demand-policies?return_records=true" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"log_path\": \"/report\", \"name\": \"on-demand-policy\", \"scan_paths\": [ \"/voll/cifs/\" ], \"scope\": { \"include_extensions\": [ \"mp*\" ], \"scan_without_extension\": true }}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "vs1"
      },
      "name": "on-demand-policy",
      "scan_paths": [
        "voll/cifs/"
      ],
      "log_path": "/report",
      "scope": {
        "max_file_size": 10737418240,
        "include_extensions": [
          "vmdk",
          "mp*"
        ],
        "scan_without_extension": true
      }
    }
  ]
}

'''

=== Updating a Vscan On-Demand policy

```

The policy being modified is identified by the UUID of the SVM and the policy name.

The API:

/api/protocols/vscan/{svm.uuid}/on-demand-policies/{name}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-demand-policies/on-demand-policy" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"schedule\": { \"name\": \"weekly\" }, \"scope\": { \"exclude_extensions\": [ \"mp3\" ], \"exclude_paths\": [ \"/vol/\" ], \"include_extensions\": [ \"vmdk\", \"mp3\" ], \"scan_without_extension\": true }}"
```

'''

=== Deleting a Vscan On-Demand policy

The policy to be deleted is identified by the UUID of the SVM and the policy name.

The API:

/api/protocols/vscan/{svm.uuid}/on-demand-policies/{name}

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/protocols/vscan/86fbc414-f140-11e8-8e22-0050568e0945/on-demand-policies/on-demand-policy" -H "accept: application/hal+json"
```

'''

[[ID221643bb0dfbf8f1adc717175aec8047]]

= Retrieve a Vscan On-Demand policy

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/vscan/{svm.uuid}/on-demand-policies`#

Introduced In: 9.6

Retrieves the Vscan On-Demand policy.

== Related ONTAP commands

* `vserver vscan on-demand-task show`

== Learn more

* xref:{relative_path}protocols_vscan_svm.uuid_on-demand-policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-demand-policies]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|scan_paths

|string

|query

|False

a|Filter by scan_paths

|log_path

|string

|query

|False

a|Filter by log_path

|scope.include_extensions

|string

|query

|False

a|Filter by scope.include_extensions

```
|scope.scan_without_extension  
|boolean  
|query  
|False  
a|Filter by scope.scan_without_extension
```

```
|scope.max_file_size  
|integer  
|query  
|False  
a|Filter by scope.max_file_size
```

```
|scope.exclude_paths  
|string  
|query  
|False  
a|Filter by scope.exclude_paths
```

```
|scope.exclude_extensions  
|string  
|query  
|False  
a|Filter by scope.exclude_extensions
```

```
|schedule.name  
|string  
|query  
|False  
a|Filter by schedule.name
```

```
|schedule.uuid  
|string  
|query  
|False  
a|Filter by schedule.uuid
```

```
|name  
|string  
|query  
|False  
a|Filter by name
```



```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1
* Max value: 120
* Min value: 0

```

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#vscan_on_demand[vscan_on_demand]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  }
}

```

```

    }
  },
  "records": [
    {
      "log_path": "/vol0/report_dir",
      "name": "task-1",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scope": {
        "exclude_extensions": [
          "mp3",
          "mp4"
        ],
        "exclude_paths": [
          "/vol1/cold-files/",
          "/vol1/cifs/names"
        ],
        "include_extensions": [
          "vmdk",
          "mp*"
        ],
        "max_file_size": "10737418240"
      }
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#schedule]
[.api-collapsible-fifth-title]

```

schedule

Schedule of the task.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|Job schedule name
```

```
|uuid
```

```
|string
```

```
a|Job schedule UUID
```

```
|===
```

```
[#scope]
```

```
[.api-collapsible-fifth-title]
```

```
scope
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|exclude_extensions
```

```
|array[string]
```

```
a|List of file extensions for which scanning is not performed.
```

```
|exclude_paths
```

```
|array[string]
```

```
a|List of file paths for which scanning must not be performed.
```

```
|include_extensions
|array[string]
a|List of file extensions to be scanned.
```

```
|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.
```

```
|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.
```

```
|===
```

```
[#vscan_on_demand]
[.api-collapsible-fifth-title]
vscan_on_demand
```

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|log_path
|string
a|The path from the Vserver root where the task report is created.
```

```
|name
|string
a|On-Demand task name
```

```
|scan_paths
|array[string]
a|List of paths that need to be scanned.
```

```
|schedule
|link:#schedule[schedule]
a|Schedule of the task.
```

```
|scope
|link:#scope[scope]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```



```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[ID61b7192ab6daa7520ffc0a8e86f09180]]
```

```
= Create a Vscan On-Demand policy
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-  
block]#`/protocols/vscan/{svm.uuid}/on-demand-policies`#
```

```
*Introduced In:* 9.6
```

Creates a Vscan On-Demand policy. Created only on a data SVM.

Important notes:

*** Only one policy can be scheduled at a time on an SVM. Use schedule name or schedule uuid to schedule an On-Demand policy.

*** Scanning must be enabled on the SVM before the policy is scheduled to run.

*** The `exclude_extensions` setting overrides the `include_extensions` setting. Set `scan_without_extension` to true to scan files without

extensions.

== Required properties

- * `svm.uuid` - Existing SVM in which to create the Vscan On-Demand policy.
- * `name` - Name of the Vscan On-Demand policy. Maximum length is 256 characters.
- * `log_path` - Path from the Vserver root where the On-Demand policy report is created.
- * `scan_paths` - List of paths that need to be scanned.

== Recommended optional properties

- * `schedule` - Scan schedule. It is recommended to set the schedule property, as it dictates when to scan for viruses.

== Default property values

If not specified in POST, the following default property values are assigned:

- * `include_extensions` - `_*`
- * `max_file_size` - `_10737418240_`
- * `scan_without_extension` - `_true_`

== Related ONTAP commands

- * ``vserver vscan on-demand-task create``
- * ``vserver vscan on-demand-task schedule``

== Learn more

* `xref:{relative_path}protocols_vscan_svm.uuid_on-demand-policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-demand-policies]`

== Parameters

[cols=5*,options=header]
|===

Name	Type	In	Required	Description
------	------	----	----------	-------------

```
|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.
```

* Default value:

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|log_path
|string
a|The path from the Vserver root where the task report is created.
```

```
|name
|string
a|On-Demand task name
```

```
|scan_paths
|array[string]
a|List of paths that need to be scanned.
```

```
|schedule
|link:#schedule[schedule]
a|Schedule of the task.
```

```
|scope
|link:#scope[scope]
a|
```

```
|===
```

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "log_path": "/vol0/report_dir",
  "name": "task-1",
  "scan_paths": [
    "/vol1/",
    "/vol2/cifs/"
  ],
  "schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": {
    "exclude_extensions": [
      "mp3",
      "mp4"
    ],
    "exclude_paths": [
      "/vol1/cold-files/",
      "/vol1/cifs/names"
    ],
    "include_extensions": [
      "vmdk",
      "mp*"
    ],
    "max_file_size": "10737418240"
  }
}
```

====

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#vscan_on_demand[vscan_on_demand]]
a|

|===


.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "log_path": "/vol0/report_dir",
      "name": "task-1",
      "scan_paths": [
        "/vol1/",
        "/vol2/cifs/"
      ],
      "schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "scope": {
        "exclude_extensions": [
          "mp3",
          "mp4"
        ],
        "exclude_paths": [
          "/vol1/cold-files/",
          "/vol1/cifs/names"
        ],
        "include_extensions": [
          "vmdk",

```

```

        "mp*"
    ],
    "max_file_size": "10737418240"
}
}
]
}
====

```

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
10027101	The file size must be in the range 1KB to 1TB
10027107	The include extensions list cannot be empty. Specify at least one extension for inclusion.
10027164	An On-Demand policy cannot be scheduled, as the Vscan is disabled. Enable the Vscan and retry the operation.
10027167	The specified schedule does not exist. Create the schedule or create a policy without specifying the schedule.
10027168	The specified scan path does not exist. The scan path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use "/" not "\")
10027169	The specified scan path is not supported for scanning.
10027173	The new On-Demand policy cannot be created as the SVM has reached the maximum number of On-Demand policies allowed. Delete an existing policy in order to create a new policy.

```

| 10027174
| The specified exclude path is invalid. The path must be specified from
the root of the SVM, and must begin with UNIX path delimiters (use "/" not
"\")

| 10027175
| An On-Demand policy cannot be scheduled as the Vserver is not in an
operational state.

| 10027176
| The log-path specified does not exist. The log path must be specified
from the root of the SVM, and must begin with UNIX path delimiters (use
"/" not "\").

| 10027177
| The log path specified is not supported.

| 10027253
| The number of paths specified exceeds the configured maximum number of
paths. You cannot specify more than the maximum number of configured
paths.

| 10027254
| The number of extensions specified exceeds the configured maximum number
of extensions. You cannot specify more than the maximum number of
configured extensions.

| 10027255
| Another policy is already scheduled. Only one policy per SVM is allowed
to be scheduled at any one time. Create a policy without specifying a
schedule.
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#schedule]
[.api-collapsible-fifth-title]
schedule

Schedule of the task.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Job schedule name

|uuid
|string
a|Job schedule UUID

|===

[#scope]
[.api-collapsible-fifth-title]
scope

[cols=3*,options=header]
|===
|Name

```



```

|Type
|Description

|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions
|array[string]
a|List of file extensions to be scanned.

|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.

|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.

|===

[#vscan_on_demand]
[.api-collapsible-fifth-title]
vscan_on_demand

Use On-Demand scanning to check files for viruses on a schedule. An On-
Demand policy defines the scope of an On-Demand scan.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|log_path
|string

```

a|The path from the Vserver root where the task report is created.

|name

|string

a|On-Demand task name

|scan_paths

|array[string]

a|List of paths that need to be scanned.

|schedule

|link:#schedule[schedule]

a|Schedule of the task.

|scope

|link:#scope[scope]

a|

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDb30ed4bf4489151cb48d051a536d0379]]
= Delete a Vscan On-Demand configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/vscan/{svm.uuid}/on-demand-policies/{name}`#

*Introduced In:* 9.6

```

Deletes the Vscan On-Demand configuration.

== Related ONTAP commands

* `vserver vscan on-demand-task delete`

== Learn more

* [xref:{relative_path}protocols_vscan_svm.uuid_on-demand-policies_endpoint_overview.html\[DOC /protocols/vscan/{svm.uuid}/on-demand-policies\]](#)

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|name

|string

|path

|True

a|

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|===

== Response

Status: 200, Ok

== Error

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

```

```
//end collapsible .Definitions block
====
```

```
[[IDfd4adbcce8d03da3ac215a27900cc705]]
```

= Retrieve the Vscan On-Demand configuration for an SVM

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/protocols/vscan/{svm.uuid}/on-demand-policies/{name}`#
```

Introduced In: 9.6

Retrieves the Vscan On-Demand configuration of an SVM.

== Related ONTAP commands

* `vserver vscan on-demand-task show`

== Learn more

* xref:{relative_path}protocols_vscan_svm.uuid_on-demand-
policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-demand-
policies]

== Parameters

```
[cols=5*,options=header]  
|==
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|name  
|string  
|path  
|True  
a|
```

```
|svm.uuid  
|string
```

```

|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|log_path
|string
a|The path from the Vserver root where the task report is created.

|name
|string
a|On-Demand task name

|scan_paths
|array[string]
a|List of paths that need to be scanned.

|schedule
|link:#schedule[schedule]
a|Schedule of the task.

|scope
|link:#scope[scope]
a|

```


|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "log_path": "/vol0/report_dir",
  "name": "task-1",
  "scan_paths": [
    "/vol1/",
    "/vol2/cifs/"
  ],
  "schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": {
    "exclude_extensions": [
      "mp3",
      "mp4"
    ],
    "exclude_paths": [
      "/vol1/cold-files/",
      "/vol1/cifs/names"
    ],
    "include_extensions": [
      "vmdk",
      "mp*"
    ],
    "max_file_size": "10737418240"
  }
}
```

====

== Error

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#schedule]
[.api-collapsible-fifth-title]
schedule

Schedule of the task.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Job schedule name

```

```
|uuid
|string
a|Job schedule UUID
```

```
|===
```

```
[#scope]
[.api-collapsible-fifth-title]
scope
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.
```

```
|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.
```

```
|include_extensions
|array[string]
a|List of file extensions to be scanned.
```

```
|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.
```

```
|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID2b58ad1c0862c826b0ccecab03203072]]
= Update the Vscan On-Demand policy configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/vscan/{svm.uuid}/on-demand-policies/{name}`#
```

Introduced In: 9.6

Updates the Vscan On-Demand policy configuration of an SVM. Use schedule name or schedule UUID to schedule an On-Demand scan.

== Related ONTAP commands

```
* `vserver vscan on-demand-task modify`
* `vserver vscan on-demand-task schedule`
* `vserver vscan on-demand-task unschedule`
```

== Learn more

```
* xref:{relative_path}protocols_vscan_svm.uuid_on-demand-
policies_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/on-demand-
policies]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
```

```

|Required
|Description

|name
|string
|path
|True
a|

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

```

```

|===

```

```

== Request Body

```

```

[cols=3*,options=header]

```

```

|===

```

```

|Name
|Type
|Description

```

```

|log_path
|string
a|The path from the Vserver root where the task report is created.

```

```

|scan_paths
|array[string]
a|List of paths that need to be scanned.

```

```

|schedule
|link:#schedule[schedule]
a|Schedule of the task.

```

```

|scope
|link:#scope[scope]
a|

```

```

|===

```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "log_path": "/vol0/report_dir",
  "scan_paths": [
    "/vol1/",
    "/vol2/cifs/"
  ],
  "schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": {
    "exclude_extensions": [
      "mp3",
      "mp4"
    ],
    "exclude_paths": [
      "/vol1/cold-files/",
      "/vol1/cifs/names"
    ],
    "include_extensions": [
      "vmdk",
      "mp*"
    ],
    "max_file_size": "10737418240"
  }
}
====
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|====
```


Error Code	Description
10027101	The file size must be in the range 1KB to 1TB
10027107	The include extensions list cannot be empty. Specify at least one extension for inclusion.
10027164	An On-Demand policy cannot be scheduled, as the Vscan is disabled. Enable the Vscan and retry the operation.
10027167	The specified schedule does not exist. Create the schedule or create a policy without specifying the schedule.
10027168	The specified scan path does not exist. The scan path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use "/" not "\")
10027169	The specified scan path is not supported for scanning.
10027174	The specified exclude path is invalid. The path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use "/" not "\")
10027175	An On-Demand policy cannot be scheduled as the SVM is not in an operational state.
10027176	The log-path specified does not exist. The log path must be specified from the root of the SVM, and must begin with UNIX path delimiters (use "/" not "\")
10027177	The log path specified is not supported.
10027253	The number of paths specified exceeds the configured maximum number of paths. You cannot specify more than the maximum number of configured paths.

```

| 10027254
| The number of extensions specified exceeds the configured maximum number
of extensions. You cannot specify more than the maximum number of
configured extensions.

| 10027255
| Another policy is already scheduled. Only one policy per SVM is allowed
to be scheduled at any one time. Update a policy without specifying a
schedule.
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#schedule]
[.api-collapsible-fifth-title]
schedule

Schedule of the task.

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|name
|string
a|Job schedule name

|uuid
|string
a|Job schedule UUID

|===

[#scope]
[.api-collapsible-fifth-title]
scope

[cols=3*,options=header]
|===
|Name
|Type
|Description

|exclude_extensions
|array[string]
a|List of file extensions for which scanning is not performed.

|exclude_paths
|array[string]
a|List of file paths for which scanning must not be performed.

|include_extensions
|array[string]
a|List of file extensions to be scanned.

|max_file_size
|integer
a|Maximum file size, in bytes, allowed for scanning.

```

```
|scan_without_extension
|boolean
a|Specifies whether or not files without any extension can be scanned.
```

```
|===
```

```
[#vscan_on_demand]
[.api-collapsible-fifth-title]
vscan_on_demand
```

Use On-Demand scanning to check files for viruses on a schedule. An On-Demand policy defines the scope of an On-Demand scan.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|log_path
|string
a|The path from the Vserver root where the task report is created.
```

```
|scan_paths
|array[string]
a|List of paths that need to be scanned.
```

```
|schedule
|link:#schedule[schedule]
a|Schedule of the task.
```

```
|scope
|link:#scope[scope]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
```

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

:leveloffset: -1

= Manage Vscan scanner-pool configuration

:leveloffset: +1

[[ID64244d2676e219fe853d9f38bd14127a]]
= Protocols Vscan svm.uuid scanner-pools endpoint overview

== Overview

A scanner-pool defines the Vscan servers and privileged users that can
connect to SVMs and a scanner policy or role determines whether a scanner-
pool is active. You can configure a scanner-pool to be used on the local
cluster or any other cluster in an MCC/DR setup.

== Examples

=== Retrieving all fields for all scanner-pools of an SVM

-----

# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/<svm-uuid>/scanner-
pools?fields=*&return_records=true&return_timeout=15" -H "accept:
application/json"

# The response:

```

```

{
  "records": [
    {
      "svm": {
        "uuid": "0e2f7c91-f227-11e8-9601-0050568ecc06"
      },
      "name": "scanner-1",
      "servers": [
        "1.1.1.1",
        "10.72.204.27"
      ],
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ],
      "role": "primary"
    },
    {
      "svm": {
        "uuid": "0e2f7c91-f227-11e8-9601-0050568ecc06"
      },
      "name": "scanner-2",
      "servers": [
        "1.1.1.1",
        "10.72.204.27"
      ],
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ],
      "role": "secondary"
    }
  ],
  "num_records": 2
}
-----

```

=== Retrieving all scanner-pools with `_role_` set as `_secondary_`

The API:

`/api/protocols/vscan/{svm.uuid}/scanner-pools`

The call:

`curl -X GET "https://<mgmt-ip>/api/protocols/vscan/<svm-uuid>/scanner-`

```
pools?role=secondary&fields=*&return_records=true&return_timeout=15" -H
"accept: application/json"
```

```
# The response:
```

```
{
  "records": [
    {
      "svm": {
        "uuid": "0e2f7c91-f227-11e8-9601-0050568ecc06",
        "name": "vs1"
      },
      "name": "scanner-2",
      "servers": [
        "1.1.1.1",
        "10.72.204.27"
      ],
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ],
      "role": "secondary",
      "cluster": {
        "uuid": "0933f9b5-f226-11e8-9601-0050568ecc06",
        "name": "Cluster3"
      }
    }
  ],
  "num_records": 1
}
----
```

```
=== Retrieving the specified scanner-pool associated with an SVM
```

```
----
```

```
# The API:
```

```
/api/protocols/vscan/{svm.uuid}/scanner-pools/{name}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/vscan/0e2f7c91-f227-11e8-
9601-0050568ecc06/scanner-pools/scanner-1?fields=*" -H "accept:
application/json"
```

```
# The response:
```

```
{
  "svm": {
```



```

    "uuid": "0e2f7c91-f227-11e8-9601-0050568ecc06",
    "name": "vs1"
  },
  "name": "scanner-1",
  "servers": [
    "1.1.1.1",
    "10.72.204.27"
  ],
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
  "role": "primary",
  "cluster": {
    "uuid": "0933f9b5-f226-11e8-9601-0050568ecc06",
    "name": "Cluster3"
  }
}

```

=== Creating a scanner-pool for an SVM with all fields specified

The API:

/api/protocols/vscan/{svm.uuid}/scanner-pools/

The call:

```

curl -X POST "https://<mgmt-ip>/api/protocols/vscan/b103be27-17b8-11e9-
b451-0050568ecd85/scanner-pools?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"cluster\": {
\"name\": \"Cluster1\", \"uuid\": \"ab746d77-17b7-11e9-b450-0050568ecd85\"
}, \"name\": \"test-scanner\", \"privileged_users\": [ \"cifs\\\\u1\",
\"cifs\\\\u2\" ], \"role\": \"primary\", \"servers\": [ \"1.1.1.1\",
\"10.72.204.27\" ] }"

```

The response:

```

{
  "num_records": 1,
  "records": [
    {
      "name": "test-scanner",
      "servers": [
        "1.1.1.1",
        "10.72.204.27"
      ],

```

```

    "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
    ],
    "role": "primary",
    "cluster": {
        "uuid": "ab746d77-17b7-11e9-b450-0050568ecd85",
        "name": "Cluster1"
    }
}
]
}
-----

```

=== Creating a scanner-pool for an SVM with an unspecified role and cluster

```

-----

# The API:
/api/protocols/vscan/{svm.uuid}/scanner-pools/

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/vscan/b103be27-17b8-11e9-b451-0050568ecd85/scanner-pools" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"name\": \"test-scanner-1\", \"privileged_users\": [ \"cifs\\\\u1\", \"cifs\\\\u2\" ], \"servers\": [ \"1.1.1.1\", \"10.72.204.27\" ]}"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "name": "test-scanner-1",
      "servers": [
        "1.1.1.1",
        "10.72.204.27"
      ],
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ]
    }
  ]
}

```

=== Updating a scanner-pool for an SVM with all of the fields specified

The API:

/api/protocols/vscan/{svm.uuid}/scanner-pools/{name}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/0e2f7c91-f227-11e8-9601-0050568ecc06/scanner-pools/test-scanner-1" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"cluster\": { \"name\": \"Cluster3\", \"uuid\": \"0933f9b5-f226-11e8-9601-0050568ecc06\" }, \"privileged_users\": [ \"cifs\\\\\\\\u1\", \"cifs\\\\\\\\u2\" ], \"role\": \"secondary\", \"servers\": [ \"1.1.1.1\", \"10.72.204.27\" ] }"
```

=== Updating the "role" of a scanner-pool for an SVM

The API:

/api/protocols/vscan/{svm.uuid}/scanner-pools/{name}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/protocols/vscan/0e2f7c91-f227-11e8-9601-0050568ecc06/scanner-pools/test-scanner-1" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"cluster\": { \"name\": \"Cluster3\", \"uuid\": \"0933f9b5-f226-11e8-9601-0050568ecc06\" }, \"role\": \"primary\" }"
```

=== Deleting a scanner-pool for a specified SVM

The API:

/api/protocols/vscan/{svm.uuid}/scanner-pools/{name}

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/protocols/vscan/0e2f7c91-f227-11e8-9601-0050568ecc06/scanner-pools/test-scanner-1" -H "accept: application/json"
```

```
[[ID66e9a81c2fca720f7eecc43486904635]]
```

= Retrieve the Vscan scanner-pool configuration for an SVM

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/protocols/vscan/{svm.uuid}/scanner-pools`#
```

Introduced In: 9.6

Retrieves the Vscan scanner-pool configuration of an SVM.

== Related ONTAP commands

```
* `vserver vscan scanner-pool show`  
* `vserver vscan scanner-pool privileged-users show`  
* `vserver vscan scanner-pool servers show`  
* `vserver vscan scanner-pool show-active`
```

== Learn more

```
* xref:{relative_path}protocols_vscan_svm.uuid_scanner-  
pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-  
pools]
```

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|cluster.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by cluster.name
```

```
|cluster.uuid
```

```
|string
```

```
|query  
|False  
a|Filter by cluster.uuid
```

```
|role  
|string  
|query  
|False  
a|Filter by role
```

```
|privileged_users  
|string  
|query  
|False  
a|Filter by privileged_users
```

```
|name  
|string  
|query  
|False  
a|Filter by name
```

```
|servers  
|string  
|query  
|False  
a|Filter by servers
```

```
|svm.uuid  
|string  
|path  
|True  
a|UUID of the SVM to which this object belongs.
```

```
|fields  
|array[string]  
|query  
|False  
a|Specify the fields to return.
```

```

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#vscan_scanner_pool[vscan_scanner_pool]]
a|

|===

```

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "scanner-1",
      "privileged_users": [

```

```

        "cifs\\u1",
        "cifs\\u2"
    ],
    "role": "string",
    "servers": [
        "1.1.1.1",
        "10.72.204.27",
        "vmwin204-27.fsct.nb"
    ]
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",

```



```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#cluster_reference]
[.api-collapsible-fifth-title]
cluster_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#vscan_scanner_pool]
[.api-collapsible-fifth-title]

```

vscan_scanner_pool

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|cluster
```

```
|link:#cluster_reference[cluster_reference]
```

```
a|
```

```
|name
```

```
|string
```

a|Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".".

```
|privileged_users
```

```
|array[string]
```

a|Specifies a list of privileged users. A valid form of privileged user-name is "domain-name\user-name". Privileged user-names are stored and treated as case-insensitive strings. Virus scanners must use one of the registered privileged users for connecting to clustered Data ONTAP for exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.

```
* example: ["cifs\u1", "cifs\u2"]
```

```
* Introduced in: 9.6
```

```
* x-omitempty: true
```

```
|role
```

```
|string
```

a|Specifies the role of the scanner pool. The possible values are:

```
*** primary    - Always active.
```

```
*** secondary - Active only when none of the primary external virus-scanning servers are connected.
```

*** idle - Always inactive.

|servers
|array[string]
a|Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.

* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]
* Introduced in: 9.6
* x-omitempty: true

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

[[IDfa43a9fa7554ce24d0f30deb5f17b4f0]]
= Create a Vscan scanner-pool configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/vscan/{svm.uuid}/scanner-pools`#

*Introduced In:* 9.6

Creates a Vscan scanner-pool configuration for a specified SVM. You can
create a scanner-pool with all fields specified or only mandatory fields
specified.

Important notes:

* A scanner-pool must have servers and privileged users specified.
* If the role or cluster is not specified, the scanner-pool is created on
the local cluster with the role set as primary.

```

*`Only one of the fields cluster-uuid or cluster-name is required.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the Vscan configuration.

* `name` - Scanner-pool name.

* `privileged_users` - List of privileged users.

* `servers` - List of server IP addresses or FQDNs.

== Recommended optional properties

* `role` - Setting a role for a scanner-pool is recommended.

* `cluster` - Passing the cluster name or UUID (or both) in a multi-cluster environment is recommended.

== Default property values

If not specified in POST, the following default property values are assigned:

* `role` - `_primary_`

* `cluster.name` - Local cluster name.

* `cluster.uuid` - Local cluster UUID.

== Related ONTAP commands

* ``vserver vscan scanner-pool create``

* ``vserver vscan scanner-pool apply-policy``

* ``vserver vscan scanner-pool privileged-users add``

* ``vserver vscan scanner-pool servers add``

== Learn more

* `xref:{relative_path}protocols_vscan_svm.uuid_scanner-pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-pools]`

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

```

|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster_reference[cluster_reference]
a|

|name
|string
a|Specifies the name of the scanner pool. Scanner pool name can be up to
256 characters long and is a string that can only contain any combination
of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".".

|privileged_users
|array[string]
a|Specifies a list of privileged users. A valid form of privileged user-
name is "domain-name\user-name". Privileged user-names are stored and
treated as case-insensitive strings. Virus scanners must use one of the
registered privileged users for connecting to clustered Data ONTAP for

```

exchanging virus-scanning protocol messages and to access file for scanning, remediating and quarantining operations.

```
* example: ["cifs\u1", "cifs\u2"]
* Introduced in: 9.6
* x-omitempty: true
```

|role

|string

a|Specifies the role of the scanner pool. The possible values are:

*** primary - Always active.

*** secondary - Active only when none of the primary external virus-scanning servers are connected.

*** idle - Always inactive.

|servers

|array[string]

a|Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.

```
* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]
* Introduced in: 9.6
* x-omitempty: true
```

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "scanner-1",
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
}
```



```
"role": "string",  
"servers": [  
  "1.1.1.1",  
  "10.72.204.27",  
  "vmwin204-27.fsct.nb"  
]  
}  
====
```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#vscan_scanner_pool[vscan_scanner_pool]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "scanner-1",
      "privileged_users": [
        "cifs\\u1",
        "cifs\\u2"
      ],
      "role": "string",
      "servers": [
        "1.1.1.1",
        "10.72.204.27",
        "vmwin204-27.fsct.nb"
      ]
    }
  ]
}
=====

== Error

```

ONTAP Error Response Codes

```
//start table
[cols=2*,options=header]
|===
//header
| Error Code | Description
//end header
//end row
//start row
|10027086 +
//end row
//start row
|The specified list of servers contain one or more entries that cannot be
resolved
//end row
//start row
|10027258 +
//end row
//start row
|The specified cluster_name does not exist
//end row
//start row
|10027256 +
//end row
//start row
|The specified cluster_uuid does not exist
//end row
//start row
|10027257 +
//end row
//start row
|The specified cluster_name and cluster_uuid are valid but belong to
different clusters
//end row
//start row
|10027248 +
//end row
//start row
|Scanner-pool created successfully but failed to activate
//end row
//start row
|10027107 +
//end row
//start row
```

```

|The list of privileged users or list of servers specified is empty
//end row
//start row
|10027108 +
//end row
//start row
|The list of privileged users specified contains an invalid entry
//end row
//start row
|10027063 +
//end row
//start row
|Attempting to modify a scanner-pool on an administrative SVM with a data
SVM
//end row
|===
//end table

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```
[#cluster_reference]
[.api-collapsible-fifth-title]
cluster_reference
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#vscan_scanner_pool]
[.api-collapsible-fifth-title]
vscan_scanner_pool
```

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|cluster
|link:#cluster_reference[cluster_reference]
a|
```

```
|name
|string
a|Specifies the name of the scanner pool. Scanner pool name can be up to 256 characters long and is a string that can only contain any combination of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".".
```

```

|privileged_users
|array[string]
a|Specifies a list of privileged users. A valid form of privileged user-
name is "domain-name\user-name". Privileged user-names are stored and
treated as case-insensitive strings. Virus scanners must use one of the
registered privileged users for connecting to clustered Data ONTAP for
exchanging virus-scanning protocol messages and to access file for
scanning, remedying and quarantining operations.

* example: ["cifs\u1", "cifs\u2"]
* Introduced in: 9.6
* x-omitempty: true


|role
|string
a|Specifies the role of the scanner pool. The possible values are:

*** primary    - Always active.

*** secondary - Active only when none of the primary external virus-
scanning servers are connected.

*** idle       - Always inactive.


|servers
|array[string]
a|Specifies a list of IP addresses or FQDN for each Vscan server host
names which are allowed to connect to clustered ONTAP.

* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]
* Introduced in: 9.6
* x-omitempty: true


|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

```
//end collapsible .Definitions block
```

```
====
```

```
[[ID4454b1b615fb5afdddb47d80873cd1d9]]
```

```
= Delete the Vscan scanner-pool configuration
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/protocols/vscan/{svm.uuid}/scanner-pools/{name}`#
```

```
*Introduced In:* 9.6
```

Deletes a Vscan scanner-pool configuration.

Important notes:

* The Vscan scanner-pool DELETE endpoint deletes all of the Vscan scanner-pools for a specified SVM.

* If a Vscan is enabled, it requires at least one scanner-pool to be in the active state. Therefore, disable Vscan on the specified SVM so all the scanner-pools configured on that SVM can be deleted.

```
== Related ONTAP commands
```

```
* `vserver vscan scanner-pool delete`
```

```
== Learn more
```

```
* xref:{relative_path}protocols_vscan_svm.uuid_scanner-  
pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-  
pools]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|name
```



```

|string
|path
|True
a|

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error Code | Description
//end header
//end row
//start row
|10027070 +
//end row
//start row
|Attempting to delete a scanner-pool but it is the only active scanner-
pool for a Vscan enabled on the SVM
//end row
//start row
|10027064 +
//end row
//start row
|Attempting to delete a scanner-pool with a data SVM which was created
with an administrative SVM
//end row
|===
//end table

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

|===

//end collapsible .Definitions block

====

[[IDaa110cac7624f666c64d00b551917af2]]

= Retrieve the Vscan scanner-pool configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/vscan/{svm.uuid}/scanner-pools/{name}`#

Introduced In: 9.6

Retrieves the configuration of a specified scanner-pool of an SVM.

== Related ONTAP commands

- * `vserver vscan scanner-pool show`
- * `vserver vscan scanner-pool privileged-users show`
- * `vserver vscan scanner-pool servers show`
- * `vserver vscan scanner-pool show-active`

== Learn more

* xref:{relative_path}protocols_vscan_svm.uuid_scanner-pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-pools]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|name

|string

```

|path
|True
a|

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster_reference[cluster_reference]
a|

|name
|string
a|Specifies the name of the scanner pool. Scanner pool name can be up to
256 characters long and is a string that can only contain any combination
of ASCII-range alphanumeric characters a-z, A-Z, 0-9), "_", "-" and ".".

|privileged_users
|array[string]
a|Specifies a list of privileged users. A valid form of privileged user-
name is "domain-name\user-name". Privileged user-names are stored and
treated as case-insensitive strings. Virus scanners must use one of the
registered privileged users for connecting to clustered Data ONTAP for

```

exchanging virus-scanning protocol messages and to access file for scanning, remedying and quarantining operations.

* example: ["cifs\u1", "cifs\u2"]

* Introduced in: 9.6

* x-omitempty: true

|role

|string

a|Specifies the role of the scanner pool. The possible values are:

*** primary - Always active.

*** secondary - Active only when none of the primary external virus-scanning servers are connected.

*** idle - Always inactive.

|servers

|array[string]

a|Specifies a list of IP addresses or FQDN for each Vscan server host names which are allowed to connect to clustered ONTAP.

* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]

* Introduced in: 9.6

* x-omitempty: true

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
```

```

"name": "scanner-1",
"privileged_users": [
  "cifs\\u1",
  "cifs\\u2"
],
"role": "string",
"servers": [
  "1.1.1.1",
  "10.72.204.27",
  "vmwin204-27.fsct.nb"
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",

```

```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#cluster_reference]

```


[.api-collapsible-fifth-title]

cluster_reference

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|

|uuid

|string

a|

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDe3ae498cf2e835ae0322b5e00bbe2844]]
= Update the Vscan scanner-pool configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/vscan/{svm.uuid}/scanner-pools/{name}`#

*Introduced In:* 9.6

Updates the Vscan scanner-pool configuration of an SVM.

```

Important notes:

- * Along with servers and privileged-users, the role of a scanner-pool can also be updated with the cluster on which a scanner-pool is allowed.
- * If role is specified and cluster isn't, then role is applied to the local cluster.

== Related ONTAP commands

```
* `vserver vscan scanner-pool modify`  
* `vserver vscan scanner-pool apply-policy`  
* `vserver vscan scanner-pool privileged-users add`  
* `vserver vscan scanner-pool privileged-users remove`  
* `vserver vscan scanner-pool servers remove`  
* `vserver vscan scanner-pool servers add`
```

== Learn more

* xref:{relative_path}protocols_vscan_svm.uuid_scanner-pools_endpoint_overview.html[DOC /protocols/vscan/{svm.uuid}/scanner-pools]

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|name  
|string  
|path  
|True  
a|
```

```
|svm.uuid  
|string  
|path  
|True  
a|UUID of the SVM to which this object belongs.
```

```

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster_reference[cluster_reference]
a|

|privileged_users
|array[string]
a|Specifies a list of privileged users. A valid form of privileged user-
name is "domain-name\user-name". Privileged user-names are stored and
treated as case-insensitive strings. Virus scanners must use one of the
registered privileged users for connecting to clustered Data ONTAP for
exchanging virus-scanning protocol messages and to access file for
scanning, remediating and quarantining operations.

* example: ["cifs\u1", "cifs\u2"]
* Introduced in: 9.6
* x-omitempty: true

|role
|string
a|Specifies the role of the scanner pool. The possible values are:

*** primary    - Always active.

*** secondary - Active only when none of the primary external virus-
scanning servers are connected.

*** idle       - Always inactive.

|servers
|array[string]
a|Specifies a list of IP addresses or FQDN for each Vscan server host
names which are allowed to connect to clustered ONTAP.

```

```
* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]
* Introduced in: 9.6
* x-omitempty: true
```

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "cluster": {
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "privileged_users": [
    "cifs\\u1",
    "cifs\\u2"
  ],
  "role": "string",
  "servers": [
    "1.1.1.1",
    "10.72.204.27",
    "vmwin204-27.fsct.nb"
  ]
}
```

====

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

//start table

[cols=2*,options=header]

|===

//header

| Error Code | Description

//end header

```

//end row
//start row
|10027258 +
//end row
//start row
|The specified cluster_name does not exist
//end row
//start row
|10027256 +
//end row
//start row
|The specified cluster_uuid does not exist
//end row
//start row
|10027257 +
//end row
//start row
|The specified cluster_name and cluster_uuid are valid but belong to
different clusters
//end row
//start row
|10027248 +
//end row
//start row
|Scanner-pool updated successfully but failed to apply the specified role
//end row
//start row
|10027107 +
//end row
//start row
|The list of privileged users or list of servers specified is empty
//end row
//start row
|10027108 +
//end row
//start row
|The list of privileged users specified contains an invalid entry
//end row
//start row
|10027063 +
//end row
//start row
|Attempting to modify a scanner-pool on an administrative SVM with a data
SVM
//end row
|===

```

```

//end table

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#cluster_reference]
[.api-collapsible-fifth-title]
cluster_reference

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

```

```
|===
```

```
[#vscan_scanner_pool]  
[.api-collapsible-fifth-title]  
vscan_scanner_pool
```

Scanner pool is a set of attributes which are used to validate and manage connections between clustered ONTAP and external virus-scanning server, or "Vscan server".

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description  
  
|cluster  
|link:#cluster_reference[cluster_reference]  
a|
```

```
|privileged_users  
|array[string]  
a|Specifies a list of privileged users. A valid form of privileged user-  
name is "domain-name\user-name". Privileged user-names are stored and  
treated as case-insensitive strings. Virus scanners must use one of the  
registered privileged users for connecting to clustered Data ONTAP for  
exchanging virus-scanning protocol messages and to access file for  
scanning, remedying and quarantining operations.
```

```
* example: ["cifs\u1", "cifs\u2"]  
* Introduced in: 9.6  
* x-omitempty: true
```

```
|role  
|string  
a|Specifies the role of the scanner pool. The possible values are:
```

```
*** primary    - Always active.
```

```
*** secondary - Active only when none of the primary external virus-  
scanning servers are connected.
```

```
*** idle       - Always inactive.
```



```
|servers
|array[string]
a|Specifies a list of IP addresses or FQDN for each Vscan server host
names which are allowed to connect to clustered ONTAP.
```

```
* example: ["1.1.1.1", "10.72.204.27", "vmwin204-27.fsct.nb"]
* Introduced in: 9.6
* x-omitempty: true
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
:leveloffset: -1
```

```
= NDMP
```

```
:leveloffset: +1
```

```
[[IDf7f10bbdf0dcf6043540fdbde42e6704]]
= NDMP overview
```

NDMP (Network Data Management Protocol) provides an open standard for network-based backup of network-attached storage (NAS). NDMP minimizes the

coding required for different applications by providing standard commands for backing up and restoring file servers. NDMP increases the speed and efficiency of NAS data protection, as data can bypass backup servers and be written directly to tape storage. Starting from ONTAP 8.2, the software supports SVM aware NDMP backups. This helps in optimizing NDMP backup performance by choosing efficient data transfer paths and is fully compatible with integrated non-disruptive operations and volume mobility capabilities of ONTAP software. In order to support this feature, a CAB extension must be implemented by backup vendors. These APIs allow you to manage the following endpoints:

- * ndmp
- * ndmp nodes
- * ndmp svms
- * ndmp sessions

== APIs

=== NDMP

These APIs are used to manage NDMP mode, SVM-scope or node-scope.

=== NDMP nodes

These APIs are used to manage node configurations.

=== NDMP SVMs

These APIs are used to manage NDMP configurations of SVMs.

=== NDMP sessions

These APIs are used to manage diagnostic information on NDMP sessions belonging to a specific SVM in the case of SVM-scope or to a specific node in the case of node-scope.

= Manage NDMP mode

:leveloffset: +1

[[IDe2ed1c2fd4d3a255483e87d7b2cc9be9]]

= Protocols NDMP endpoint overview

You can use this API to manage NDMP mode: SVM-scope or node-scope.

=== Examples

Updates NDMP mode to SVM:

```
----  
PATCH "/api/protocols/ndmp" '{"mode":"svm"}'  
----
```

Updates NDMP mode to node:

```
----  
PATCH "/api/protocols/ndmp" '{"mode":"node"}'  
----
```

[[ID518b5840d8686902e0b50be0742cd47f]]

= Retrieve the current NDMP mode

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/ndmp`#

Introduced In: 9.7

Retrieves the current NDMP mode.

== Related ONTAP commands

* `system services ndmp node-scope-mode status`

== Learn more

* xref:{relative_path}protocols_ndmp_endpoint_overview.html[DOC /protocols/ndmp]

== Parameters

[cols=5*,options=header]

|===

|Name
|Type
|In
|Required
|Description

|mode
|string
|query
|False
a|Filter by mode

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction.

Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|mode
|string
a|Indicates whether NDMP is in node-scoped or SVM-scoped mode.

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "mode": "string"
}
=====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

```

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

====

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#href]

[.api-collapsible-fifth-title]

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

|string

a|

|===


```

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

[[ID72bc42f796608a5ca8b84ca5a3bc6c4c]]
= Update the NDMP mode

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/ndmp`#

*Introduced In:* 9.7

Updates the NDMP mode.

== Related ONTAP commands

* `system services ndmp node-scope-mode`

== Learn more

```

```

* xref:{relative_path}protocols_ndmp_endpoint_overview.html[DOC
/protocols/ndmp]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mode
|string
a|Indicates whether NDMP is in node-scoped or SVM-scoped mode.

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]

```

```
{
  "mode": "string"
}
====
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|mode
|string
a|Indicates whether NDMP is in node-scoped or SVM-scoped mode.

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "mode": "string"
}
====

== Error
```

Status: Default, Error

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
```

```

href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#cluster_ndmp_properties]
[.api-collapsible-fifth-title]
cluster_ndmp_properties

[cols=3*,options=header]
|===
|Name
|Type
|Description

|mode
|string
a|Indicates whether NDMP is in node-scoped or SVM-scoped mode.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code

```

```

|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
:leveloffset: -1
```

= Manage NDMP node configurations

```
:leveloffset: +1
```

```
[[ID2f044f59b27e70c4d700ec9bc3abf1eb]]
```

= Protocols NDMP nodes endpoint overview

You can use this API to manage node configurations. The authentication types "plaintext" and "plaintext_sso" are used to show that the password uses clear text. Also, they contain no differences for NDMP node scope.

=== Examples

Updates "enabled" and "authentication_types" fields:

```
----
```

```
PATCH "/api/protocols/ndmp/nodes/13bb2092-458b-11e9-9c06-0050568ea64e"
'{"enabled":"false","authentication_types":["plaintext"]}'
```

```
----
```

Updates the "user" field:

```
----
```

```
PATCH "/api/protocols/ndmp/nodes/13bb2092-458b-11e9-9c06-0050568ea64e"
'{"user":"user22"}'
```

```
----
```

```
[[ID8f0caea8cdd89bf8529337fc84036881]]
```

= Retrieve NDMP node configurations for all nodes

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/protocols/ndmp/nodes`#
```

Introduced In: 9.7

Retrieves NDMP node configurations for all of the nodes.

== Related ONTAP commands

* ``system services ndmp show``

== Learn more

* `xref:{relative_path}protocols_ndmp_nodes_endpoint_overview.html` [DOC
/protocols/ndmp/nodes]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|authentication_types

|string

|query

|False

a|Filter by authentication_types

|user

|string

|query

|False

a|Filter by user

|enabled

|boolean

|query

|False

a|Filter by enabled

|node.name

|string


```
|query
|False
a|Filter by node.name
```

```
|node.uuid
|string
|query
|False
a|Filter by node.uuid
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
 |_links
|link:#_links[_links]
a|
```

```
 |num_records
|integer
a|Number of Records
```

```
 |records
|array[link:#ndmp_node[ndmp_node]]
a|
```

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourceLink"
```

```

    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_types": [
        "plaintext",
        "challenge"
      ],
      "enabled": 1,
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "user": "ndmp_user"
    }
  ]
}
====

== Error

```

Status: Default

ONTAP Error Response codes

|===

Error code	Description
------------	-------------

68812801	
----------	--

	Node-scoped operations are not allowed in an SVM-scope.
--	---------------------------------------------------------

68812804	
----------	--

| Failed to get the node name from the specified node UUID.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "error": {

 "arguments": [

 {

 "code": "string",

 "message": "string"

 }

],

 "code": "4",

 "message": "entry doesn't exist",

 "target": "uuid"

 }

}

====

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#href]

[.api-collapsible-fifth-title]

```

href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self

```

```

|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

Storage system node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#ndmp_node]
[.api-collapsible-fifth-title]
ndmp_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean
a|Is the NDMP service enabled?

|node
|link:#node[node]
a|Storage system node

|user
|string
a|NDMP user ID

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID712327fedb0a24b344d19f60a6113467]]
= Retrieve the NDMP node configuration for a node

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/ndmp/nodes/{node.uuid}`#

*Introduced In:* 9.7

Retrieves an NDMP node configuration for a specific node.

```


== Related ONTAP commands

* ``system services ndmp show``

== Learn more

* `xref:{relative_path}protocols_ndmp_nodes_endpoint_overview.html` [DOC
/protocols/ndmp/nodes]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|node.uuid

|string

|path

|True

a|Node UUID

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

```
|_links
|link:#_links[_links]
a|
```

```
|authentication_types
|array[string]
a|NDMP authentication types.
```

```
|enabled
|boolean
a|Is the NDMP service enabled?
```

```
|node
|link:#node[node]
a|Storage system node
```

```
|user
|string
a|NDMP user ID
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```

```

    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"user": "ndmp_user"
}
====

== Error

```

Status: Default

ONTAP Error Response codes

```

|====
| Error code | Description

| 68812801
| Node-scoped operations are not allowed in an SVM-scope.

| 68812802
| The UUID is not valid.

| 68812804
| Failed to get the node name from the specified node UUID.
|====

```

[cols=3*,options=header]

```

|====
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```

|====

```

```

.Example error
[%collapsible%closed]
====

```

```

[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

Storage system node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

```
//end collapsible .Definitions block
====

[[IDdffcd4cba66ec53ef94f549f4adbc0ac]]
= Update the NDMP node configuration for a node

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/ndmp/nodes/{node.uuid}`#

*Introduced In:* 9.7

Updates the NDMP node configuration for a specific node.

== Related ONTAP commands

* `system services ndmp modify`

== Learn more

* xref:{relative_path}protocols_ndmp_nodes_endpoint_overview.html[DOC
/protocols/ndmp/nodes]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|node.uuid
|string
|path
|True
a|Node UUID

|===

== Request Body
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean
a|Is the NDMP service enabled?

|node
|link:#node[node]
a|Storage system node

|password
|string
a|NDMP password. This can only be set and cannot be read back.

|user
|string
a|NDMP user ID

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "node": {

```



```
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "password": "string",
  "user": "ndmp_user"
}
====

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean
a|Is the NDMP service enabled?

|node
|link:#node[node]
a|Storage system node

|password
|string
a|NDMP password. This can only be set and cannot be read back.

|user
|string
a|NDMP user ID

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "password": "string",
  "user": "ndmp_user"
}
====

== Error
```

Status: Default

ONTAP Error Response codes

```
|===
| Error code | Description
|
| 68812800
| The user is required to enable NDMP on a node.
|
| 68812801
| Node-scoped operations are not allowed in an SVM-scope.
|
| 68812802
| The UUID is not valid.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

Storage system node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#ndmp_node]
[.api-collapsible-fifth-title]
ndmp_node

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean
a|Is the NDMP service enabled?

|node
|link:#node[node]
a|Storage system node

|password
|string
a|NDMP password. This can only be set and cannot be read back.

|user
|string
a|NDMP user ID

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string

```

```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
:leveloffset: -1
```

```
= Manage NDMP session configurations
```

```
:leveloffset: +1
```

```
[[IDc416de8118893425c273a7dc1d87dc58]]
```

```
= Protocols NDMP sessions endpoint overview
```

You can use this API to manage diagnostic information on NDMP sessions belonging to a specific SVM in the case of SVM-scope or to a specific node in the case of node-scope.

=== Examples

Retrieves NDMP session details under node-scope:

```
----
```

```
GET "/api/protocols/ndmp/sessions/9b372ce7-3a4b-11e9-a7f8-0050568e3d73/2000"
```

```
----
```

Retrieves NDMP session details under SVM-scope:

```
----
```

```
GET "/api/protocols/ndmp/sessions/13bb2092-458b-11e9-9c06-0050568ea604/2000:4000"
```

```
----
```

Deletes NDMP session details under node-scope:

```
----
```

```
DELETE "/api/protocols/ndmp/sessions/9b372ce7-3a4b-11e9-a7f8-0050568e3d73/2000"
```

```
----
```

Deletes NDMP session details under SVM-scope:

```
----
```

```
DELETE "/api/protocols/ndmp/sessions/13bb2092-458b-11e9-9c06-0050568ea604/2000:4000"
```

[[ID1f7436ea17c3b0f44016b4643366ffdd]]

= Retrieve NDMP sessions

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/ndmp/sessions`#

Introduced In: 9.7

Retrieves a collection of NDMP sessions. In the case of SVM-scope, if this API is executed on a data IP, it displays the list of NDMP sessions under the specified SVM; otherwise it displays the list of NDMP sessions for all the SVMs under the cluster. In the case of node-scope, it displays the list of NDMP sessions for all nodes.

== Related ONTAP commands

* `vserver services ndmp probe`

* `system services ndmp probe`

== Learn more

* xref:{relative_path}protocols_ndmp_sessions_endpoint_overview.html[DOC /protocols/ndmp/sessions]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|tape_mode

|string

|query

|False

a|Filter by tape_mode

```
|backup_engine
|string
|query
|False
a|Filter by backup_engine
```

```
|scsi.target_id
|integer
|query
|False
a|Filter by scsi.target_id
```

```
|scsi.lun_id
|integer
|query
|False
a|Filter by scsi.lun_id
```

```
|scsi.device_id
|string
|query
|False
a|Filter by scsi.device_id
```

```
|scsi.host_adapter
|integer
|query
|False
a|Filter by scsi.host_adapter
```

```
|client_address
|string
|query
|False
a|Filter by client_address
```

```
|client_port
|integer
|query
```



```
|False  
a|Filter by client_port
```

```
|svm.uuid  
|string  
|query  
|False  
a|Filter by svm.uuid
```

```
|svm.name  
|string  
|query  
|False  
a|Filter by svm.name
```

```
|data.state  
|string  
|query  
|False  
a|Filter by data.state
```

```
|data.operation  
|string  
|query  
|False  
a|Filter by data.operation
```

```
|data.connection.address  
|string  
|query  
|False  
a|Filter by data.connection.address
```

```
|data.connection.port  
|integer  
|query  
|False  
a|Filter by data.connection.port
```

```
|data.connection.type
```

```
|string
|query
|False
a|Filter by data.connection.type
```

```
|data.reason
|string
|query
|False
a|Filter by data.reason
```

```
|data.bytes_processed
|integer
|query
|False
a|Filter by data.bytes_processed
```

```
|tape_device
|string
|query
|False
a|Filter by tape_device
```

```
|node.name
|string
|query
|False
a|Filter by node.name
```

```
|node.uuid
|string
|query
|False
a|Filter by node.uuid
```

```
|mover.state
|string
|query
|False
a|Filter by mover.state
```

```
|mover.bytes_moved
|integer
|query
|False
a|Filter by mover.bytes_moved
```

```
|mover.connection.address
|string
|query
|False
a|Filter by mover.connection.address
```

```
|mover.connection.port
|integer
|query
|False
a|Filter by mover.connection.port
```

```
|mover.connection.type
|string
|query
|False
a|Filter by mover.connection.type
```

```
|mover.reason
|string
|query
|False
a|Filter by mover.reason
```

```
|mover.mode
|string
|query
|False
a|Filter by mover.mode
```

```
|source_address
|string
|query
|False
```

a|Filter by source_address

|data_path

|string

|query

|False

a|Filter by data_path

|id

|string

|query

|False

a|Filter by id

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|return_records

|boolean

|query

|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP

returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of Records.

|records

|array[link:#ndmp_session[ndmp_session]]

a|

|===

.Example response

[%collapsible%closed]

====

```
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "backup_engine": "string",
      "client_address": "string",
      "data": {
        "bytes_processed": "5000",
        "connection": {
          "address": "string",
          "port": "18600",
          "type": "local"
        },
        "operation": "backup",
        "reason": "aborted",
        "state": "halted"
      },
      "data_path": "/vserver1/vol1",
      "id": "string",
      "mover": {
        "bytes_moved": "645120",
        "connection": {
          "address": "string",
          "port": "18600",
          "type": "local"
        },
        "mode": "read",
        "reason": "end_of_media",
        "state": "connected"
      },
      "node": {
        "_links": {
          "self": {

```

```

        "href": "/api/resourcelink"
      },
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scsi": {
    "device_id": "string"
  },
  "source_address": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tape_device": "nrst0a",
  "tape_mode": "write"
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response codes

|===

| Error code | Description

| 68812805

| Failed to obtain the NDMP mode of the operation.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

```
|error
|link:#error[error]
a|
```

```
|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
```



```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#ndmp_connect]
[.api-collapsible-fifth-title]
ndmp_connect

```

Indicates the NDMP connection attributes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|address
```

```
|string
```

a|Indicates the NDMP data connection address.

```
|port
```

```
|integer
```

a|Indicates the NDMP data connection port.

```
|type
```

```
|string
```

a|The different NDMP connection types.

```
|===
```

```
[#ndmp_data]
```

```
[.api-collapsible-fifth-title]
```

ndmp_data

Information about the NDMP data server.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|bytes_processed
```

```
|integer
```

a|Indicates the NDMP data bytes processed.

```
|connection
```

```
|link:#ndmp_connect[ndmp_connect]
```

a|Indicates the NDMP connection attributes.

|operation

|string

a|Indicates the NDMP data server operation.

|reason

|string

a|If the NDMP mover or data server is in the paused or halted state, then this property provides additional information for the state.

|state

|string

a|The various states of the NDMP state machine during a backup or restore workflow.

|===

[#ndmp_mover]

[.api-collapsible-fifth-title]

ndmp_mover

Information about the NDMP mover.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|bytes_moved

|integer

a|Indicates the NDMP mover bytes moved.

|connection

|link:#ndmp_connect[ndmp_connect]

a|Indicates the NDMP connection attributes.

|mode

```

|string
a|The modes of a NDMP mover instance.

|reason
|string
a|If the NDMP mover or data server is in the paused or halted state, then
this property provides additional information for the state.

|state
|string
a|The various states of the NDMP state machine during a backup or restore
workflow.

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#ndmp_scsi]
[.api-collapsible-fifth-title]
ndmp_scsi

```

Information about the NDMP SCSI server.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|device_id
```

```
|string
```

```
a|Indicates the NDMP SCSI device ID.
```

```
|host_adapter
```

```
|integer
```

```
a|Indicates the NDMP SCSI host adapter.
```

```
|lun_id
```

```
|integer
```

```
a|Indicates the NDMP SCSI LUN ID.
```

```
|target_id
```

```
|integer
```

```
a|Indicates the NDMP SCSI target ID.
```

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#ndmp_session]
[.api-collapsible-fifth-title]
ndmp_session

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|backup_engine
|string
a|Indicates the NDMP backup engine.

|client_address
|string
a|Indicates the NDMP client address.

|client_port
|integer
a|Indicates the NDMP client port.

|data
|link:#ndmp_data[ndmp_data]
a|Information about the NDMP data server.

```

```

|data_path
|string
a|Indicates the NDMP backup or restore path.

|id
|string
a|NDMP session identifier.

|mover
|link:#ndmp_mover[ndmp_mover]
a|Information about the NDMP mover.

|node
|link:#node[node]
a|

|scsi
|link:#ndmp_scsi[ndmp_scsi]
a|Information about the NDMP SCSI server.

|source_address
|string
a|Indicates the NDMP local address on which connection was established.

|svm
|link:#svm[svm]
a|

|tape_device
|string
a|Indicates the NDMP tape device.

|tape_mode
|string
a|The modes of a NDMP mover instance.

|===

[#error_arguments]

```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message


```

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDe6cec39de036e5ed67581cce4349de6f]]
= Delete an NDMP session

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/ndmp/sessions/{owner.uuid}/{session.id}`#

*Introduced In:* 9.7

Deletes a specific NDMP session.

== Related ONTAP commands

* `vserver services ndmp kill`
* `system services ndmp kill`

== Learn more

* xref:{relative_path}protocols_ndmp_sessions_endpoint_overview.html[DOC
/protocols/ndmp/sessions]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string

```

```
|path
|True
a|NDMP node UUID

|session.id
|string
|path
|True
a|Session identifier

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response codes

```
|===
| Error code | Description

| 68812802
| The UUID is not valid.

| 68812803
| Failed to get the SVM name from the specified SVM UUID.

| 68812804
| Failed to get the node name from the specified node UUID.

| 68812805
| Failed to obtain the NDMP mode of operation.

| 68812806
| UUID and Session ID are required.
|===
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

```
//end collapsible .Definitions block
====

[[IDe02fda4ca7c18dc0cd54522f2720f288]]
= Retrieve NDMP session details

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/ndmp/sessions/{owner.uuid}/{session.id}`#

*Introduced In:* 9.7

Retrieves the details of a specific NDMP session.

== Related ONTAP commands

* `vserver services ndmp probe`
* `system services ndmp probe`

== Learn more

* xref:{relative_path}protocols_ndmp_sessions_endpoint_overview.html[DOC
/proTOCOLS/NDMP/SESSIONS]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|owner.uuid
|string
|path
|True
a|The NDMP node or SVM UUID based on whether NDMP is operating in node-
scope or SVM-scope mode.

|session.id
```

```
|string
|path
|True
a|NDMP session identifier

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|backup_engine
|string
a|Indicates the NDMP backup engine.

|client_address
|string
a|Indicates the NDMP client address.

|client_port
|integer
a|Indicates the NDMP client port.

|data
|link:#ndmp_data[ndmp_data]
a|Information about the NDMP data server.
```

```

|data_path
|string
a|Indicates the NDMP backup or restore path.

|id
|string
a|NDMP session identifier.

|mover
|link:#ndmp_mover[ndmp_mover]
a|Information about the NDMP mover.

|node
|link:#node[node]
a|

|scsi
|link:#ndmp_scsi[ndmp_scsi]
a|Information about the NDMP SCSI server.

|source_address
|string
a|Indicates the NDMP local address on which connection was established.

|svm
|link:#svm[svm]
a|

|tape_device
|string
a|Indicates the NDMP tape device.

|tape_mode
|string
a|The modes of a NDMP mover instance.

|===

```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "backup_engine": "string",
  "client_address": "string",
  "data": {
    "bytes_processed": "5000",
    "connection": {
      "address": "string",
      "port": "18600",
      "type": "local"
    },
    "operation": "backup",
    "reason": "aborted",
    "state": "halted"
  },
  "data_path": "/vserver1/vol1",
  "id": "string",
  "mover": {
    "bytes_moved": "645120",
    "connection": {
      "address": "string",
      "port": "18600",
      "type": "local"
    },
    "mode": "read",
    "reason": "end_of_media",
    "state": "connected"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
}
```



```

"scsi": {
  "device_id": "string"
},
"source_address": "string",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tape_device": "nrst0a",
"tape_mode": "write"
}
====

== Error

```

Status: Default

ONTAP Error Response codes

```

|===
| Error code | Description
| 68812802
| The UUID is not valid.
|===

```

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#ndmp_connect]
```

```
[.api-collapsible-fifth-title]
```

```
ndmp_connect
```

Indicates the NDMP connection attributes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|address
```

```
|string
```

a|Indicates the NDMP data connection address.

```
|port
```

```
|integer
```

a|Indicates the NDMP data connection port.

```
|type
```

```
|string
```

a|The different NDMP connection types.

```
|===
```

```
[#ndmp_data]
```

```
[.api-collapsible-fifth-title]
```

```
ndmp_data
```

Information about the NDMP data server.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|bytes_processed
```

```
|integer
```

```
a|Indicates the NDMP data bytes processed.
```

```
|connection
```

```
|link:#ndmp_connect[ndmp_connect]
```

```
a|Indicates the NDMP connection attributes.
```

```
|operation
```

```
|string
```

```
a|Indicates the NDMP data server operation.
```

```
|reason
```

```
|string
```

```
a|If the NDMP mover or data server is in the paused or halted state, then  
this property provides additional information for the state.
```

```
|state
```

```
|string
```

```
a|The various states of the NDMP state machine during a backup or restore  
workflow.
```

```
|===
```

```
[#ndmp_mover]
```

```
[.api-collapsible-fifth-title]
```

```
ndmp_mover
```

Information about the NDMP mover.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|bytes_moved
|integer
a|Indicates the NDMP mover bytes moved.

|connection
|link:#ndmp_connect[ndmp_connect]
a|Indicates the NDMP connection attributes.

|mode
|string
a|The modes of a NDMP mover instance.

|reason
|string
a|If the NDMP mover or data server is in the paused or halted state, then
this property provides additional information for the state.

|state
|string
a|The various states of the NDMP state machine during a backup or restore
workflow.

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links

```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#ndmp_scsi]
```

```
[.api-collapsible-fifth-title]
```

```
ndmp_scsi
```

Information about the NDMP SCSI server.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|device_id
```

```
|string
```

```
a|Indicates the NDMP SCSI device ID.
```

```
|host_adapter
```

```
|integer
```

```
a|Indicates the NDMP SCSI host adapter.
```

```
|lun_id
```

```
|integer
```

```
a|Indicates the NDMP SCSI LUN ID.
```

```
|target_id
```

```
|integer
```

```
a|Indicates the NDMP SCSI target ID.
```

```

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string

```

```

a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

```


= Manage NDMP configuration of SVM

:leveloffset: +1

[[IDa7a65cbc61b5a4d6aaf93d39012923bf]]

= Protocols NDMP svms endpoint overview

You can use this API to manage NDMP configurations of SVMs.

=== Examples

Updates the "enabled" field:

```
PATCH "/api/protocols/ndmp/svms/9b372ce7-3a4b-11e9-a7f8-0050568e3d73"
'{"enabled":"false"}'
```

Updates the "authentication_types" field:

```
PATCH "/api/protocols/ndmp/svms/9b372ce7-3a4b-11e9-a7f8-0050568e3d73"
'{"authentication_types":["challenge"]}'
```

[[ID20fbb39bccc63b4426cff7e067108646]]

= Retrieve the NDMP configurations for all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/ndmp/svms`#

Introduced In: 9.7

Retrieves NDMP configurations for all SVMs.

== Related ONTAP commands

* `vserver services ndmp show`

== Learn more

```
* xref:{relative_path}protocols_ndmp_svms_endpoint_overview.html[DOC
/protocols/ndmp/svms]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|authentication_types
|string
|query
|False
a|Filter by authentication_types
```

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|fields
|array[string]
|query
```

```

|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of Records


|records
|array[link:#ndmp_svm[ndmp_svm]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_types": [
        "plaintext",
```

```

        "challenge"
    ],
    "enabled": 1,
    "svm": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ]
  }
}

```

```

    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self

```

```

|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

```

|===

[#ndmp_svm]
[.api-collapsible-fifth-title]
ndmp_svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean
a|Is the NDMP service enabled on the SVM?

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```



```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

[[ID3a632109cec7433a5ba2af87de339238]]

= Retrieve the NDMP configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/ndmp/svms/{svm.uuid}`#

Introduced In: 9.7

Retrieves an NDMP configuration for a specific SVM.

== Related ONTAP commands

* `vserver services ndmp show`

== Learn more

* xref:{relative_path}protocols_ndmp_svms_endpoint_overview.html[DOC /protocols/ndmp/svms]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|SVM UUID

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean
a|Is the NDMP service enabled on the SVM?

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
```

```

"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

```

}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID2b4d8c4c6f4c3019847a8aeb437f8104]]
= Update the NDMP configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/ndmp/svms/{svm.uuid}`#

*Introduced In:* 9.7

Updates the NDMP configuration for a specific SVM.

```

== Related ONTAP commands

* ``vserver services ndmp modify``

== Learn more

* `xref:{relative_path}protocols_ndmp_svms_endpoint_overview.html` [DOC
/protocols/ndmp/svms]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|SVM UUID

|===

== Request Body

[cols=3*,options=header]
|===

|Name
|Type
|Description

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean


```
a|Is the NDMP service enabled on the SVM?
```

```
|svm  
|link:#svm[svm]  
a|
```

```
|===
```

```
.Example request
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{  
  "authentication_types": [  
    "plaintext",  
    "challenge"  
  ],  
  "enabled": 1,  
  "svm": {  
    "name": "svm1",  
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"  
  }  
}
```

```
====
```

```
== Response
```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean
a|Is the NDMP service enabled on the SVM?

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "authentication_types": [
    "plaintext",
    "challenge"
  ],
  "enabled": 1,
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error

```

Status: Default, Error

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

```

```

[#ndmp_svm]
[.api-collapsible-fifth-title]
ndmp_svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication_types
|array[string]
a|NDMP authentication types.

|enabled
|boolean
a|Is the NDMP service enabled on the SVM?

|svm
|link:#svm[svm]
a|

|===

//end collapsible .Definitions block
====

:leveloffset: -1

:leveloffset: -1

= NVMe

:leveloffset: +1

[[ID953fb152b5e7bee0963158e55c26636a]]
= NVMe overview

```

== Overview

The Non-Volatile Memory Express (NVMe) API endpoints and objects provide for configuration, provisioning and management of the NVMe-related objects.

NVMe over Fabrics (NVMe-oF) refers to the extensions and changes to the base NVMe command set to support NVMe commands over a fabric interconnect and from multiple hosts simultaneously. ONTAP implements elements of both NVMe and NVMe-oF. Throughout this documentation, NVMe is generally used to refer to both NVMe and NVMe-oF.

=== Fibre Channel Logins

Fibre Channel logins represent connections, formed by Fibre Channel initiators, that have successfully logged in to ONTAP. This represents the Fibre Channel login on which higher-level protocols such as Fibre Channel Protocol (FCP) and Non-Volatile Memory Express over Fibre Channel (NVMe over FC) rely.

The Fibre Channel logins REST API provides information about active Fibre Channel logins.

=== NVMe Interfaces

NVMe interfaces are network interfaces configured to support an NVMe over Fabrics protocol. The NVMe interfaces are Fibre Channel interfaces supporting an NVMe-oF data protocol. Regardless of the underlying physical and data protocol, NVMe interfaces are treated equally for the host-side application configuration. This endpoint provides a consolidated view of all NVMe interfaces for the purpose of configuring host-side applications.

The NVMe interfaces REST API provides NVMe-specific information about network interfaces configured to support an NVMe-oF protocol.

==== Learn More

* `_Fibre Channel Interfaces_` found in the `_networking_` section. Fibre Channel interfaces are the logical endpoints for Fibre Channel network connections to an SVM.

=== NVMe Services

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service

for a given SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

The Non-Volatile Memory Express (NVMe) service REST API allows you to create, update, delete, and discover NVMe services for SVMs.

=== NVMe Subsystem Controllers

Non-Volatile Memory Express (NVMe) subsystem controllers represent dynamic connections between hosts and a storage solution.

The NVMe subsystem controllers REST API provides information about connected hosts.

=== NVMe Subsystem Maps

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

The NVMe subsystem map REST API allows you to create, delete, and discover NVMe subsystem maps.

=== NVMe Subsystems

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

The NVMe subsystem REST API allows you to create, update, delete, and discover NVMe subsystems. It also allows you to add and remove NVMe hosts that can access the subsystem and associated namespaces.

=== NVMe Namespaces

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

The NVMe namespace REST API allows you to create, update, delete, and discover NVMe namespaces.

```
= View NVMe interfaces
```

```
:leveloffset: +1
```

```
[[IDfe7e64141652dc76b9b20cb6a25be1c5]]
```

```
= Protocols NVMe interfaces endpoint overview
```

```
== Overview
```

NVMe interfaces are network interfaces configured to support an NVMe over Fabrics (NVMe-oF) protocol. The NVMe interfaces are Fibre Channel (FC) interfaces supporting an NVMe-oF data protocol. Regardless of the underlying physical and data protocol, NVMe interfaces are treated equally for host-side application configuration. This endpoint provides a consolidated view of all NVMe interfaces for the purpose of configuring host-side applications.

The NVMe interfaces REST API provides NVMe-specific information about network interfaces configured to support an NVMe-oF protocol.

NVMe interfaces must be created using the protocol-specific endpoints for FC interfaces. See `xref:{relative_path}post-network-fc-interfaces.html[POST /network/fc/interfaces]` . After creation, the interfaces are available via this interface.

```
== Examples
```

```
=== Retrieving summary information for all NVMe interfaces
```

```
----
```

```
# The API:
```

```
GET /api/protocols/nvme/interfaces
```

```
# The call:
```

```
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/interfaces' -H 'accept: application/hal+json'
```

```
# The response:
```

```
{
  "records": [
    {
      "svm": {
        "uuid": "013e2c44-0d30-11e9-a684-005056bbdb14",
```

```

    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/013e2c44-0d30-11e9-a684-005056bbdb14"
      }
    }
  },
  "uuid": "74d69872-0d30-11e9-a684-005056bbdb14",
  "name": "nvme1",
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/interfaces/74d69872-0d30-11e9-a684-005056bbdb14"
    }
  }
},
{
  "svm": {
    "uuid": "013e2c44-0d30-11e9-a684-005056bbdb14",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/013e2c44-0d30-11e9-a684-005056bbdb14"
      }
    }
  },
  "uuid": "77ded991-0d30-11e9-a684-005056bbdb14",
  "name": "nvme2",
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/interfaces/77ded991-0d30-11e9-a684-005056bbdb14"
    }
  }
}
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/nvme/interfaces"
  }
}
}
----

'''

```



```

=== Retrieving detailed information for a specific NVMe interface

----

# The API:
GET /api/protocols/nvme/interfaces/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/interfaces/77ded991-0d30-11e9-a684-005056bbdb14' -H 'accept: application/hal+json'

# The response:
{
  "svm": {
    "uuid": "013e2c44-0d30-11e9-a684-005056bbdb14",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/013e2c44-0d30-11e9-a684-005056bbdb14"
      }
    }
  },
  "uuid": "77ded991-0d30-11e9-a684-005056bbdb14",
  "name": "nvme2",
  "enabled": true,
  "node": {
    "name": "node1",
    "uuid": "cd4d47fd-0d2e-11e9-a684-005056bbdb14",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/cd4d47fd-0d2e-11e9-a684-005056bbdb14"
      }
    }
  },
  "transport_address": "nn-0x2003005056bbdb14:pn-0x2005005056bbdb14",
  "fc_interface": {
    "wwnn": "20:03:00:50:56:bb:db:14",
    "wwpn": "20:05:00:50:56:bb:db:14",
    "port": {
      "name": "1a",
      "uuid": "081ec491-0d2f-11e9-a684-005056bbdb14",
      "node": {
        "name": "node1"
      }
    },
    "_links": {

```

```

        "self": {
            "href": "/api/network/fc/ports/081ec491-0d2f-11e9-a684-005056bbdb14"
        }
    },
    "_links": {
        "self": {
            "href": "/api/network/fc/interfaces/77ded991-0d30-11e9-a684-005056bbdb14"
        }
    }
},
"_links": {
    "self": {
        "href": "/api/protocols/nvme/interfaces/77ded991-0d30-11e9-a684-005056bbdb14"
    }
}
}
}
----

'''

```

```

[[IDdf6d9fa52320b8bc5d54edaac51b7d44]]
= Retrieve NVMe interfaces

```

```

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nvme/interfaces`#

```

Introduced In: 9.6

Retrieves NVMe interfaces.

== Related ONTAP commands

* `vserver nvme show-interface`

== Learn more

* xref:{relative_path}protocols_nvme_interfaces_endpoint_overview.html[DOC /protocols/nvme/interfaces]

== Parameters

```
[cols=5*,options=header]  
|==
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|fc_interface.port.name  
|string  
|query  
|False  
a|Filter by fc_interface.port.name
```

```
|fc_interface.port.node.name  
|string  
|query  
|False  
a|Filter by fc_interface.port.node.name
```

```
|fc_interface.port.uuid  
|string  
|query  
|False  
a|Filter by fc_interface.port.uuid
```

```
|fc_interface.wwnn  
|string  
|query  
|False  
a|Filter by fc_interface.wwnn
```

```
|fc_interface.wwpn  
|string  
|query  
|False  
a|Filter by fc_interface.wwpn
```

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|transport_address
|string
|query
|False
a|Filter by transport_address
```

```
|node.name
|string
|query
|False
a|Filter by node.name
```

```
|node.uuid
|string
|query
|False
a|Filter by node.uuid
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1
* Max value: 120
* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_interface[nvme_interface]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
```

```

    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "fc_interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "port": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "0a",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
        },
        "wwnn": "20:00:00:50:56:b4:13:a9",
        "wwpn": "20:00:00:50:56:b4:13:a8"
      },
      "name": "lif1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
    },
  ],

```

```

        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "transport_address": "nn-0x200a00a0989062da:pn-0x200100a0989062da",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```



```

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]

```

`_links`

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#node]
```

```
[.api-collapsible-fifth-title]
```

```
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#port]
```

```
[.api-collapsible-fifth-title]
```

```
port
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the FC port.
```

```
|node
|link:#node[node]
a|The node on which the FC port is located.
```

```
|uuid
|string
a|The unique identifier of the FC port.
```

```
|===
```

```
[#fc_interface]
[.api-collapsible-fifth-title]
fc_interface
```

The attributes specific to a Fibre Channel-based NVMe interface.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|port
|link:#port[port]
a|An FC port is the physical port of an FC adapter on a cluster node that
can be connected to an FC network.
```

```
|wwnn
|string
a|The WWNN (world wide node name) of the Fibre Channel NVMe interface.
```

```
|wwpn
|string
a|The WWPN (world wide port name) of the Fibre Channel NVMe interface.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
```

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#nvme_interface]

[.api-collapsible-fifth-title]

nvme_interface

NVMe interfaces are network interfaces configured to support an NVMe over Fabrics (NVMe-oF) protocol. The NVMe interfaces are Fibre Channel interfaces supporting an NVMe-oF data protocol. Regardless of the underlying physical and data protocol, NVMe interfaces are treated equally for host-side application configuration. This endpoint provides a consolidated view of all NVMe interfaces for the purpose of configuring host-side applications.

NVMe interfaces must be created using the protocol-specific endpoints for Fibre Channel interfaces. See [xref:{relative_path}post-network-fc-interfaces.html](#)[POST /network/fc/interfaces] . After creation, the interfaces are available via this interface.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

```

|enabled
|boolean
a|The administrative state of the NVMe interface.

|fc_interface
|link:#fc_interface[fc_interface]
a|The attributes specific to a Fibre Channel-based NVMe interface.

|name
|string
a|The name of the NVMe interface.

|node
|link:#node[node]
a|

|svm
|link:#svm[svm]
a|

|transport_address
|string
a|The transport address of the NVMe interface.

|uuid
|string
a|The unique identifier of the NVMe interface.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code

```

```

|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[ID7e7261a935356215d24b7de75577d744]]
```

= Retrieve an NVMe interface

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/protocols/nvme/interfaces/{uuid}`#
```

Introduced In: 9.6

Retrieves an NVMe interface.

== Related ONTAP commands

* `vserver nvme show-interface`

== Learn more

* xref:{relative_path}protocols_nvme_interfaces_endpoint_overview.html[DOC
/protocols/nvme/interfaces]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|The unique identifier of the NVMe interface.
```

```
|fields
```

```
|array[string]
```

```
|query
```

```
|False
```

```
a|Specify the fields to return.
```



```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|enabled
```

```
|boolean
```

```
a|The administrative state of the NVMe interface.
```

```
|fc_interface
```

```
|link:#fc_interface[fc_interface]
```

```
a|The attributes specific to a Fibre Channel-based NVMe interface.
```

```
|name
```

```
|string
```

```
a|The name of the NVMe interface.
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|transport_address
```

```
|string
```

```
a|The transport address of the NVMe interface.
```

```
|uuid
```

```
|string
```

a|The unique identifier of the NVMe interface.

|===

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "fc_interface": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "wwnn": "20:00:00:50:56:b4:13:a9",
    "wwpn": "20:00:00:50:56:b4:13:a8"
  },
  "name": "lif1",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```

    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "transport_address": "nn-0x200a00a0989062da:pn-0x200100a0989062da",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 2621462
| The supplied SVM does not exist.
|===

```

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```

|===

```

```

.Example error
[%collapsible%closed]
====

```

```

[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description

|self
|link:#href[href]
a|
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#port]
[.api-collapsible-fifth-title]
port
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the FC port.
```

```
|node
```

```
|link:#node[node]
```

```
a|The node on which the FC port is located.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the FC port.
```

```
|===
```

```
[#fc_interface]
```

```
[.api-collapsible-fifth-title]
```

```
fc_interface
```

The attributes specific to a Fibre Channel-based NVMe interface.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|port
```

```
|link:#port[port]
```

```
a|An FC port is the physical port of an FC adapter on a cluster node that  
can be connected to an FC network.
```

```
|wwnn
```

```
|string
```

```
a|The WWNN (world wide node name) of the Fibre Channel NVMe interface.
```

```
|wwpn
|string
a|The WWPN (world wide port name) of the Fibre Channel NVMe interface.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```

a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
= Manage NVMe services
```

```
:leveloffset: +1
```

```
[[IDd6f5263598b7f707caf2af6e8e1e8c14]]
= Protocols NVMe services endpoint overview
```

```
== Overview
```

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for an SVM. An SVM's NVMe service must be created before NVMe host

initiators can connect to the SVM.

The Non-Volatile Memory Express (NVMe) service REST API allows you to create, update, delete, and discover NVMe services for SVMs.

== Performance monitoring

Performance of the SVM can be monitored by the ``metric.++`` and ``statistics.++`` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The ``metric.++`` properties denote an average whereas ``statistics.++`` properties denote a real-time monotonically increasing value aggregated across all nodes.

== Examples

=== Creating an NVMe service for an SVM

The simplest way to create an NVMe service is to specify only the SVM, either by name or UUID. By default, the new NVMe service is enabled.

In this example, the ``return_records`` query parameter is used to retrieve the new NVMe service object in the REST response.

The API:

POST /api/protocols/nvme/services

The call:

```
curl -X POST 'https://<mgmt-  
ip>/api/protocols/nvme/services?return_records=true' -H 'accept:  
application/hal+json' -d '{ "svm": { "name": "svm1" } }'
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "svm": {  
        "uuid": "bfb1be0-dc69-11e8-b29f-005056bb7341",  
        "name": "svm1",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/bfb1be0-dc69-11e8-b29f-005056bb7341"  
          }  
        }  
      }  
    }  
  ],  
}
```

```

    "enabled": true,
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-
005056bb7341"
      }
    }
  }
]
}
----

'''

```

=== Retrieving the NVMe services for all SVMs in the cluster

```

----

# The API:
GET /api/protocols/nvme/services

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/services' -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "ab60c350-dc68-11e8-9711-005056bbe408",
        "name": "svm0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/ab60c350-dc68-11e8-9711-005056bbe408"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/services/ab60c350-dc68-11e8-9711-
005056bbe408"
        }
      }
    },
    {

```

```

    "svm": {
      "uuid": "bfb1beb0-dc69-11e8-b29f-005056bb7341",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/bfb1beb0-dc69-11e8-b29f-005056bb7341"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-005056bb7341"
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/services"
    }
  }
}
-----
'''

```

=== Retrieving details for a specific NVMe service

The NVMe service is identified by the UUID of its SVM.

The API:

```
GET /api/protocols/nvme/services/{svm.uuid}
```

The call:

```
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-005056bb7341' -H 'accept: application/hal+json'
```

The response:

```

{
  "svm": {
    "uuid": "bfb1beb0-dc69-11e8-b29f-005056bb7341",
    "name": "svm1",

```

```

    "_links": {
      "self": {
        "href": "/api/svm/svms/bfb1beb0-dc69-11e8-b29f-005056bb7341"
      }
    }
  },
  "enabled": true,
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-005056bb7341"
    }
  }
}

```

'''

=== Disabling an NVMe service

Disabling an NVMe service shuts down all active NVMe connections for the SVM and prevents the creation of new NVMe connections.

The NVMe service to update is identified by the UUID of its SVM.

The API:

```
PATCH /api/protocols/nvme/services/{svm.uuid}
```

The call:

```
curl -X PATCH 'https://<mgmt-ip>/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-005056bb7341' -H 'accept: application/hal+json' -d '{"enabled": "false" }'
```

You can retrieve the NVMe service to confirm the change.

The API:

```
GET /api/protocols/nvme/services/{svm.uuid}
```

The call:

```
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-005056bb7341' -H 'accept: application/hal+json'
```

```

# The response:
{
  "svm": {
    "uuid": "bfb1beb0-dc69-11e8-b29f-005056bb7341",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/bfb1beb0-dc69-11e8-b29f-005056bb7341"
      }
    }
  },
  "enabled": false,
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/services/bfb1beb0-dc69-11e8-b29f-005056bb7341"
    }
  }
}
----

'''

=== Deleting an NVMe service

The NVMe service must be disabled before it can be deleted. In addition,
all NVMe interfaces, subsystems, and subsystem maps associated with the
SVM must first be deleted.

The NVMe service to delete is identified by the UUID of its SVM.

----

# The API:
DELETE /api/protocols/nvme/services/{svm.uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/services/bfb1beb0-
dc69-11e8-b29f-005056bb7341' -H 'accept: application/hal+json'
----

[[ID748b4d618c202d866a2326f00b319429]]
= Retrieve NVMe services

```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nvme/services`#
```

Introduced In: 9.6

Retrieves NVMe services.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `statistics.+++`

* `metric.+++`

== Related ONTAP commands

* `vserver nvme show`

== Learn more

* [xref:{relative_path}protocols_nvme_services_endpoint_overview.html\[DOC /protocols/nvme/services\]](#)

== Parameters

```
[cols=5*,options=header]
|==
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|metric.throughput.write
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by metric.throughput.write
```

* Introduced in: 9.7

```
|metric.throughput.read  
|integer  
|query  
|False  
a|Filter by metric.throughput.read
```

* Introduced in: 9.7

```
|metric.throughput.total  
|integer  
|query  
|False  
a|Filter by metric.throughput.total
```

* Introduced in: 9.7

```
|metric.duration  
|string  
|query  
|False  
a|Filter by metric.duration
```

* Introduced in: 9.7

```
|metric.timestamp  
|string  
|query  
|False  
a|Filter by metric.timestamp
```

* Introduced in: 9.7

```
|metric.status  
|string  
|query  
|False  
a|Filter by metric.status
```

* Introduced in: 9.7


```
|metric.iops.total
|integer
|query
|False
a|Filter by metric.iops.total
```

* Introduced in: 9.7

```
|metric.iops.read
|integer
|query
|False
a|Filter by metric.iops.read
```

* Introduced in: 9.7

```
|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other
```

* Introduced in: 9.7

```
|metric.iops.write
|integer
|query
|False
a|Filter by metric.iops.write
```

* Introduced in: 9.7

```
|metric.latency.total
|integer
|query
|False
a|Filter by metric.latency.total
```

* Introduced in: 9.7

```
|metric.latency.read
|integer
```

```
|query
|False
a|Filter by metric.latency.read
```

* Introduced in: 9.7

```
|metric.latency.other
|integer
|query
|False
a|Filter by metric.latency.other
```

* Introduced in: 9.7

```
|metric.latency.write
|integer
|query
|False
a|Filter by metric.latency.write
```

* Introduced in: 9.7

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|statistics.latency_raw.total
```

```
|integer
|query
|False
a|Filter by statistics.latency_raw.total
```

* Introduced in: 9.7

```
|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read
```

* Introduced in: 9.7

```
|statistics.latency_raw.other
|integer
|query
|False
a|Filter by statistics.latency_raw.other
```

* Introduced in: 9.7

```
|statistics.latency_raw.write
|integer
|query
|False
a|Filter by statistics.latency_raw.write
```

* Introduced in: 9.7

```
|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp
```

* Introduced in: 9.7

```
|statistics.iops_raw.total
|integer
|query
```

```

|False
a|Filter by statistics.iops_raw.total

* Introduced in: 9.7

|statistics.iops_raw.read
|integer
|query
|False
a|Filter by statistics.iops_raw.read

* Introduced in: 9.7

|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other

* Introduced in: 9.7

|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write

* Introduced in: 9.7

|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write

* Introduced in: 9.7

|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read

```

* Introduced in: 9.7

|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total

* Introduced in: 9.7

|statistics.status
|string
|query
|False
a|Filter by statistics.status

* Introduced in: 9.7

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout

```

|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_service[nvme_service]]
a|

|===

```

```
.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "latency": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "status": "ok",
        "throughput": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "timestamp": "2017-01-25T11:20:13Z"
      }
    }
  ]
}
```

```

    },
    "statistics": {
      "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error

```



```

|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|next  
|link:href[href]  
a|
```

```
|self  
|link:href[href]  
a|
```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|self  
|link:href[href]  
a|
```

```
|===
```

```
[#iops]  
[.api-collapsible-fifth-title]  
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

latency

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be

metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```

|===

[#metric]
[.api-collapsible-fifth-title]
metric

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in

```

value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#iops_raw]

[.api-collapsible-fifth-title]

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
```


be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw

|link:#latency_raw[latency_raw]

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

```
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_service]
[.api-collapsible-fifth-title]
nvme_service
```

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for an SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

An NVMe service is identified by the UUID of its SVM.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|enabled
|boolean
a|The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.
```

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
|link:#metric[metric]
a|
```

```
|statistics
|link:#statistics[statistics]
a|
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID77ae5cae5aaf4345a94ba40f5e50a3f8]]
= Create an NVMe service

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/nvme/services`#

*Introduced In:* 9.6

Creates an NVMe service.

== Required properties

* `svm.uuid` or `svm.name` - The existing SVM in which to create the NVMe

```

service.

== Related ONTAP commands

* `vserver nvme create`

== Learn more

* [xref:{relative_path}protocols_nvme_services_endpoint_overview.html\[DOC /protocols/nvme/services\]](#)

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
|link:#metric[metric]
a|

|statistics
|link:#statistics[statistics]
a|

|svm
|link:#svm[svm]
a|

|===
```

```
.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "metric": {
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "statistics": {
```

```

    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_service[nvme_service]]
a|

|===

```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "metric": {
        "duration": "PT15S",
        "iops": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "latency": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "status": "ok",
        "throughput": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "timestamp": "2017-01-25T11:20:13Z"
      },
      "statistics": {
        "iops_raw": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "latency_raw": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "status": "ok",
        "throughput_raw": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "timestamp": "2017-01-25T11:20:13Z"
      }
    }
  ]
}
```



```

    },
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description

| 1115127
| The cluster lacks a valid NVMe license.

| 2621462
| The supplied SVM does not exist.

| 2621507
| NVMe is not allowed for the specified SVM.

| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5374893
| The SVM is stopped. The SVM must be running to create an NVMe service.

| 72089650
| An NVMe service already exists for the specified SVM.

| 72089900
| An NVMe service cannot be creating in an SVM that is configured for a
SAN protocol.
|====

```

```
== Definitions
```

```
[.api-def-first-level]
```

```
.See Definitions
```

```
[%collapsible%closed]
```

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
```

```
[.api-collapsible-fifth-title]
```

```
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
```

```
[.api-collapsible-fifth-title]
```

```
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

```
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
```

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

```
a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.
```

```
|status
```

```
|string
```

```
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with
```


"backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_service]
```

[.api-collapsible-fifth-title]

nvme_service

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for an SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

An NVMe service is identified by the UUID of its SVM.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

|metric

|link:#metric[metric]

a|

|statistics

|link:#statistics[statistics]

a|

|svm

|link:#svm[svm]

a|

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

```
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[IDbc10c617cd3eab0d10e627c988d7d8dc]]
```

```
= Delete an NVMe service
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/protocols/nvme/services/{svm.uuid}`#
```

```
*Introduced In:* 9.6
```

Deletes an NVMe service. An NVMe service must be disabled before it can be deleted. In addition, all NVMe interfaces, subsystems, and subsystem maps associated with the SVM must first be deleted.

```
== Related ONTAP commands
```

```
* `vserver nvme delete`
```

```
== Learn more
```

```
* xref:{relative_path}protocols_nvme_services_endpoint_overview.html[DOC  
/protocols/nvme/services]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|The unique identifier of the SVM whose NVMe service is to be deleted.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 2621462
```

```
| The supplied SVM does not exist.
```

```
| 72089651
```

```
| The supplied SVM does not have an NVMe service.
```

```
| 72089653
```

```
| There are subsystems associated with the NVMe service SVM. The  
subsystems must be removed before deleting the NVMe service.
```

```
| 72089654
```

```
| There are NVMe-oF LIFs associated with the NVMe service SVM. The LIFs  
must be removed before deleting the NVMe service.
```

```
| 72090028
```

```
| The NVMe service is enabled. The NVMe service must be disabled before it  
can be deleted.
```

```
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

====

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

```

|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID0ca3b87f39673481f0c56ba4ca0eac7f]]
= Retrieve an NVMe service

```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nvme/services/{svm.uuid}`#

*Introduced In:* 9.6

Retrieves an NVMe service.

== Related ONTAP commands

* `vserver nvme show`

== Learn more

* xref:{relative_path}protocols_nvme_services_endpoint_overview.html[DOC
/protocols/nvme/services]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|The unique identifier of the SVM whose NVMe service is to be retrieved.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```



```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|The administrative state of the NVMe service. The NVMe service can be
disabled to block all NVMe connectivity to the SVM.

This is optional in POST and PATCH. The default setting is _true_
(enabled) in POST.

|metric
|link:#metric[metric]
a|

|statistics
|link:#statistics[statistics]
a|

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
```

```

"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"statistics": {
  "iops_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
  "_links": {

```

```

    "self": {
      "href": "/api/resourcelink"
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
|
| 2621462
| The supplied SVM does not exist.
|
| 72089651
| The supplied SVM does not have an NVMe service.
|===

```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```
|===
```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{

```

```

"error": {
  "arguments": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

|===

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|self

|link:#href[href]

a|

|===

[#iops]

[.api-collapsible-fifth-title]

iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#latency]

```
[.api-collapsible-fifth-title]
```

latency

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput]
```

```
[.api-collapsible-fifth-title]
```

throughput

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```

|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency

```

```
|link:#latency[latency]
```

```
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
```

```
|string
```

```
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
```

```
|throughput
```

```
|link:#throughput[throughput]
```

```
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
```

```
|string
```

```
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
```

```
[.api-collapsible-fifth-title]
```

```
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
```

```

|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[IDa7bb723ce0c3d1e35de96d6ceab686bc]]
```

= Update an NVMe service

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-  
block]#`/protocols/nvme/services/{svm.uuid}`#
```

Introduced In: 9.6

Updates an NVMe service.

== Related ONTAP commands

* `vserver nvme modify`

== Learn more

* xref:{relative_path}protocols_nvme_services_endpoint_overview.html[DOC
/protocols/nvme/services]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|The unique identifier of the SVM whose NVMe service is to be updated.
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description

|enabled
|boolean
a|The administrative state of the NVMe service. The NVMe service can be
disabled to block all NVMe connectivity to the SVM.
```

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
|link:#metric[metric]
a|

|statistics
|link:#statistics[statistics]
a|

|svm
|link:#svm[svm]
a|

|===
```

.Example request

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "metric": {
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
  },
}
```

```

    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```

|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write

```

```

|integer
a|Peformance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

[cols=3*,options=header]
|===
|Name
|Type
|Description

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
```

[.api-collapsible-fifth-title]

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#statistics]

[.api-collapsible-fifth-title]

statistics

[cols=3*,options=header]

|===

|Name

|Type

|Description

|iops_raw

|link:#iops_raw[iops_raw]

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type

```


|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#nvme_service]

[.api-collapsible-fifth-title]

nvme_service

A Non-Volatile Memory Express (NVMe) service defines the properties of the NVMe controller target for an SVM. There can be at most one NVMe service for an SVM. An SVM's NVMe service must be created before NVMe host initiators can connect to the SVM.

An NVMe service is identified by the UUID of its SVM.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|The administrative state of the NVMe service. The NVMe service can be disabled to block all NVMe connectivity to the SVM.

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

|metric

|link:#metric[metric]

a|

|statistics

```
|link:#statistics[statistics]
```

```
a|
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

```
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[ID59687bfd4d36f9723b99083149bd44a8]]
= Retrieve NVMe protocol historical performance metrics
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nvme/services/{svm.uuid}/metrics`#
```

```
*Introduced In:* 9.7
```

Retrieves historical performance metrics for NVMe protocol of an SVM.

```
== Parameters
```

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|throughput.total
```

```
|integer
|query
|False
a|Filter by throughput.total
```

```
|throughput.read
|integer
|query
|False
a|Filter by throughput.read
```

```
|throughput.other
|integer
|query
|False
a|Filter by throughput.other
```

```
|throughput.write
|integer
|query
|False
a|Filter by throughput.write
```

```
|duration
|string
|query
|False
a|Filter by duration
```

```
|timestamp
|string
|query
|False
a|Filter by timestamp
```

```
|iops.total
|integer
|query
|False
a|Filter by iops.total
```

```
|iops.read
|integer
|query
|False
a|Filter by iops.read
```

```
|iops.other
|integer
|query
|False
a|Filter by iops.other
```

```
|iops.write
|integer
|query
|False
a|Filter by iops.write
```

```
|status
|string
|query
|False
a|Filter by status
```

```
|latency.total
|integer
|query
|False
a|Filter by latency.total
```

```
|latency.read
|integer
|query
|False
a|Filter by latency.read
```

```
|latency.other
|integer
|query
|False
```

a|Filter by latency.other

|latency.write

|integer

|query

|False

a|Filter by latency.write

|svm.uuid

|string

|path

|True

a|Unique identifier of the SVM.

|interval

|string

|query

|False

a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.

The period for each time range is as follows:

* 1h: Metrics over the most recent hour sampled over 15 seconds.

* 1d: Metrics over the most recent day sampled over 5 minutes.

* 1w: Metrics over the most recent week sampled over 30 minutes.

* 1m: Metrics over the most recent month sampled over 2 hours.

* 1y: Metrics over the most recent year sampled over a day.

* Default value: 1

* enum: ["1h", "1d", "1w", "1m", "1y"]

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records
```

```
|records
|array[link:#records[records]]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
```



```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

```

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]

```

`_links`

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]  
[.api-collapsible-fifth-title]  
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|other  
|integer  
a|Performance metric for other I/O operations. Other I/O operations can be  
metadata operations, such as directory lookups and so on.
```

```
|read  
|integer  
a|Performance metric for read I/O operations.
```

```
|total  
|integer  
a|Performance metric aggregated over all types of I/O operations.
```

```
|write  
|integer  
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]  
[.api-collapsible-fifth-title]  
throughput
```

The rate of throughput bytes per second observed at the storage object.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#records]
[.api-collapsible-fifth-title]
records

Performance numbers, such as IOPS latency and throughput.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]

```

```

a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
:leveloffset: -1
```

```
= View NVMe subsystem controllers
```

```
:leveloffset: +1
```

```
[[ID8776f7a419c388108d11907411bbe833]]
```

```
= Protocols NVMe subsystem-controllers endpoint overview
```

```
== Overview
```

Non-Volatile Memory Express (NVMe) subsystem controllers represent dynamic connections between hosts and a storage solution.

The NVMe subsystem controllers REST API provides information about connected hosts.

```
== Examples
```

```
=== Retrieving the NVMe subsystem controllers for the entire system
```

```
----
```

```
# The API:
```

```
GET /api/protocols/nvme/subsystem-controllers
```

```
# The call:
```

```
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-controllers'
-H 'accept: application/hal+json'
```



```
# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
        "name": "symmcon_fcnvme_vserver_0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
          }
        }
      },
      "subsystem": {
        "uuid": "14875240-2594-11e9-abde-00a098984313",
        "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
        "_links": {
          "self": {
            "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-00a098984313"
          }
        }
      },
      "id": "0040h",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313/0040h"
        }
      }
    },
    {
      "svm": {
        "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
        "name": "symmcon_fcnvme_vserver_0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
          }
        }
      },
      "subsystem": {
        "uuid": "14875240-2594-11e9-abde-00a098984313",
        "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
        "_links": {
```

```

        "self": {
            "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-00a098984313"
        }
    },
    "id": "0041h",
    "_links": {
        "self": {
            "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313/0041h"
        }
    }
},
{
    "svm": {
        "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
        "name": "symmcon_fcnvme_vserver_0",
        "_links": {
            "self": {
                "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
            }
        }
    },
    "subsystem": {
        "uuid": "1489d0d5-2594-11e9-94c4-00a0989a1c8e",
        "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_1",
        "_links": {
            "self": {
                "href": "/api/protocols/nvme/subsystems/1489d0d5-2594-11e9-94c4-00a0989a1c8e"
            }
        }
    },
    "id": "0040h",
    "_links": {
        "self": {
            "href": "/api/protocols/nvme/subsystem-controllers/1489d0d5-2594-11e9-94c4-00a0989a1c8e/0040h"
        }
    }
},
{
    "svm": {
        "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
        "name": "symmcon_fcnvme_vserver_0",

```

```

    "_links": {
      "self": {
        "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
      }
    },
    "subsystem": {
      "uuid": "1489d0d5-2594-11e9-94c4-00a0989a1c8e",
      "name": "symmcon_symmcon_fc_nvme_vserver_0_subsystem_1",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystems/1489d0d5-2594-11e9-94c4-00a0989a1c8e"
        }
      }
    },
    "id": "0041h",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-controllers/1489d0d5-2594-11e9-94c4-00a0989a1c8e/0041h"
      }
    }
  ],
  "num_records": 4,
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystem-controllers"
    }
  }
}
----

'''

```

=== Retrieving the NVMe subsystem controllers for a specific subsystem

The API:

GET /api/protocols/nvme/subsystem-controllers

The call:

curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-controllers?subsystem.uuid=14875240-2594-11e9-abde-00a098984313' -H

```
'accept: application/hal+json'
```

```
# The response:
```

```
{
  "records": [
    {
      "svm": {
        "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
        "name": "symmcon_fcnvme_vserver_0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
          }
        }
      },
      "subsystem": {
        "uuid": "14875240-2594-11e9-abde-00a098984313",
        "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
        "_links": {
          "self": {
            "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-00a098984313"
          }
        }
      },
      "id": "0040h",
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313/0040h"
        }
      }
    },
    {
      "svm": {
        "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
        "name": "symmcon_fcnvme_vserver_0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
          }
        }
      },
      "subsystem": {
        "uuid": "14875240-2594-11e9-abde-00a098984313",
        "name": "symmcon_symmcon_fcnvme_vserver_0_subsystem_0",
```

```

    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-00a098984313"
      }
    },
    "id": "0041h",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313/0041h"
      }
    }
  },
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313"
    }
  }
}
}
}
----

'''

```

=== Retrieving a specific NVMe subsystem controller

```

----

# The API:
GET /api/protocols/nvme/subsystem-controllers/{subsystem.uuid}/{id}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313/0040h' -H 'accept: application/hal+json'

# The response:
{
  "svm": {
    "uuid": "f0f5b928-2593-11e9-94c4-00a0989a1c8e",
    "name": "symmcon_fc_nvme_vserver_0",
    "_links": {

```

```

    "self": {
      "href": "/api/svm/svms/f0f5b928-2593-11e9-94c4-00a0989a1c8e"
    }
  },
  "subsystem": {
    "uuid": "14875240-2594-11e9-abde-00a098984313",
    "name": "symmcon_symmcon_fc_nvme_vserver_0_subsystem_0",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystems/14875240-2594-11e9-abde-00a098984313"
      }
    }
  },
  "id": "0040h",
  "interface": {
    "name": "symmcon_lif_fc_nvme_symmcon_fc_nvme_vserver_0_3a_0",
    "uuid": "falc5941-2593-11e9-94c4-00a0989a1c8e",
    "transport_address": "nn-0x200400a0989a1c8d:pn-0x200500a0989a1c8d",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/interfaces/falc5941-2593-11e9-94c4-00a0989a1c8e"
      }
    }
  },
  "node": {
    "name": "ssan-8040-94a",
    "uuid": "ebf66f05-2590-11e9-abde-00a098984313",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/ebf66f05-2590-11e9-abde-00a098984313"
      }
    }
  },
  "host": {
    "transport_address": "nn-0x20000090fae00806:pn-0x10000090fae00806",
    "nqn": "nqn.2014-08.org.nvmexpress:uuid:c2846cb1-89d2-4020-a3b0-71ce907b4eef",
    "id": "b8546ca6097349e5b1558dc154fc073b"
  },
  "io_queue": {
    "count": 4,
    "depth": [
      32,

```

```

    32,
    32,
    32
  ]
},
"admin_queue": {
  "depth": 32
},
"_links": {
  "self": {
    "href": "/api/protocols/nvme/subsystem-controllers/14875240-2594-11e9-abde-00a098984313/0040h"
  }
}
}
}
-----

```

```
[[ID131e9e78d5b2cf49246f796e23466b25]]
```

```
= Retrieve NVMe subsystem controllers
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nvme/subsystem-controllers`#
```

```
*Introduced In:* 9.6
```

```
Retrieves NVMe subsystem controllers.
```

```
== Related ONTAP commands
```

```
* `vserver nvme subsystem controller show`
```

```
== Learn more
```

```
* xref:{relative_path}protocols_nvme_subsystem-controllers_endpoint_overview.html[DOC /protocols/nvme/subsystem-controllers]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
|Type
|In
|Required
|Description

|node.name
|string
|query
|False
a|Filter by node.name
```

```
|node.uuid
|string
|query
|False
a|Filter by node.uuid
```

```
|subsystem.name
|string
|query
|False
a|Filter by subsystem.name
```

```
|subsystem.uuid
|string
|query
|False
a|Filter by subsystem.uuid
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```



```
|io_queue.count
|integer
|query
|False
a|Filter by io_queue.count
```

```
|io_queue.depth
|integer
|query
|False
a|Filter by io_queue.depth
```

```
|id
|string
|query
|False
a|Filter by id
```

```
|admin_queue.depth
|integer
|query
|False
a|Filter by admin_queue.depth
```

```
|interface.uuid
|string
|query
|False
a|Filter by interface.uuid
```

```
|interface.transport_address
|string
|query
|False
a|Filter by interface.transport_address
```

```
|interface.name
|string
|query
|False
```

a|Filter by interface.name

|host.id

|string

|query

|False

a|Filter by host.id

|host.transport_address

|string

|query

|False

a|Filter by host.transport_address

|host.nqn

|string

|query

|False

a|Filter by host.nqn

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|return_records

|boolean

|query

|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```
* Default value: 1
* Max value: 120
* Min value: 0
```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_subsystem_controller[nvme_subsystem_controller]]
a|
```

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "admin_queue": {
        "depth": 0
      },
      "host": {
        "id": "b8546ca6097349e5b1558dc154fc073b",
        "nqn": "nqn.2014-08.org.nvmexpress:uuid:c2846cb1-89d2-4020-a3b0-71ce907b4eef",
        "transport_address": "nn-0x20000090fae00806:pn-0x10000090fae00806"
      },
      "id": "0040h",
      "interface": {
        "name": "lif1",
        "transport_address": "nn-0x200400a0989a1c8d:pn-0x200500a0989a1c8d",
        "uuid": "fa1c5941-2593-11e9-94c4-00a0989a1c8e"
      },
      "io_queue": {
        "count": 0,
        "depth": [
          "integer"
        ]
      },
    },
  ],
}
```

```

"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"subsystem": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

```

```

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#admin_queue]
[.api-collapsible-fifth-title]
admin_queue

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|depth
|integer
a|The depth of the admin queue for the controller.
```

```
|===
```

```
[#host]
[.api-collapsible-fifth-title]
host
```

Properties of the connected host.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|id
|string
a|The host identifier registered with the controller.
```

```
|nqn
|string
a|The NVMe qualified name of the host.
```

```
|transport_address
|string
a|The transport address of the host.
```

```
|===
```

```
[#interface]
[.api-collapsible-fifth-title]
interface
```

The logical interface through which the host is connected.


```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the logical interface.

|transport_address
|string
a|The transport address of the logical interface.

|uuid
|string
a|The unique identifier of the logical interface.
```

```
|===
```

```
[#io_queue]
[.api-collapsible-fifth-title]
io_queue
```

Properties of the I/O queues available to the controller.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of I/O queues available to the controller.

|depth
|array[integer]
a|The depths of the I/O queues.
```

```

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#subsystem]
[.api-collapsible-fifth-title]
subsystem

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the NVMe subsystem.

```

```
|uuid
|string
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_subsystem_controller]
[.api-collapsible-fifth-title]
nvme_subsystem_controller
```

A Non-Volatile Memory Express (NVMe) subsystem controller represents a connection between a host and a storage solution.

An NVMe subsystem controller is identified by the NVMe subsystem UUID and the controller ID.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|admin_queue
|link:#admin_queue[admin_queue]
a|

|host
|link:#host[host]
a|Properties of the connected host.

|id
|string
a|The identifier of the subsystem controller. This field consists of 4
zero-filled hexadecimal digits followed by an 'h'.

|interface
|link:#interface[interface]
a|The logical interface through which the host is connected.

|io_queue
|link:#io_queue[io_queue]
a|Properties of the I/O queues available to the controller.

|node
|link:#node[node]
a|

|subsystem
|link:#subsystem[subsystem]
a|

|svm
|link:#svm[svm]
a|

|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[ID2f4784e1739736f372eafd550ca3d182]]

= Retrieve an NVMe subsystem controller

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nvme/subsystem-controllers/{subsystem.uuid}/{id}`#

Introduced In: 9.6

Retrieves an NVMe subsystem controller.

== Related ONTAP commands

* `vserver nvme subsystem controller show`

== Learn more

* xref:{relative_path}protocols_nvme_subsystem-controllers_endpoint_overview.html[DOC /protocols/nvme/subsystem-controllers]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

```

|subsystem.uuid
|string
|path
|True
a|The unique identifier of the NVMe subsystem.

|id
|string
|path
|True
a|The unique identifier of the NVMe subsystem controller.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|admin_queue
|link:#admin_queue[admin_queue]
a|

|host
|link:#host[host]
a|Properties of the connected host.

```

```
|id
|string
a|The identifier of the subsystem controller. This field consists of 4
zero-filled hexadecimal digits followed by an 'h'.
```

```
|interface
|link:#interface[interface]
a|The logical interface through which the host is connected.
```

```
|io_queue
|link:#io_queue[io_queue]
a|Properties of the I/O queues available to the controller.
```

```
|node
|link:#node[node]
a|
```

```
|subsystem
|link:#subsystem[subsystem]
a|
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs+=macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "admin_queue": {
    "depth": 0
  },
  "host": {
    "id": "b8546ca6097349e5b1558dc154fc073b",
    "nqn": "nqn.2014-08.org.nvmexpress:uuid:c2846cb1-89d2-4020-a3b0-
```



```

71ce907b4eef",
  "transport_address": "nn-0x20000090fae00806:pn-0x10000090fae00806"
},
"id": "0040h",
"interface": {
  "name": "lif1",
  "transport_address": "nn-0x200400a0989a1c8d:pn-0x200500a0989a1c8d",
  "uuid": "fal5941-2593-11e9-94c4-00a0989a1c8e"
},
"io_queue": {
  "count": 0,
  "depth": [
    "integer"
  ]
},
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"subsystem": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

ONTAP Error Response Codes

|===

| Error Code | Description

| 72090001

| The supplied subsystem identifier does not exist.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

=====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#admin_queue]
[.api-collapsible-fifth-title]
admin_queue

[cols=3*,options=header]
|===
|Name

```

```
|Type
|Description

|depth
|integer
a|The depth of the admin queue for the controller.
```

```
|===
```

```
[#host]
[.api-collapsible-fifth-title]
host
```

Properties of the connected host.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|id
|string
a|The host identifier registered with the controller.
```

```
|nqn
|string
a|The NVMe qualified name of the host.
```

```
|transport_address
|string
a|The transport address of the host.
```

```
|===
```

```
[#interface]
[.api-collapsible-fifth-title]
interface
```

The logical interface through which the host is connected.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the logical interface.

|transport_address
|string
a|The transport address of the logical interface.

|uuid
|string
a|The unique identifier of the logical interface.
```

```
|===
```

```
[#io_queue]
[.api-collapsible-fifth-title]
io_queue
```

Properties of the I/O queues available to the controller.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of I/O queues available to the controller.

|depth
|array[integer]
a|The depths of the I/O queues.
```

```

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#subsystem]
[.api-collapsible-fifth-title]
subsystem

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the NVMe subsystem.

```

```
|uuid
|string
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

```


====

:leveloffset: -1

= Manage NVMe subsystem maps

:leveloffset: +1

[[ID163f02a598b914a62a1e47d170d8491e]]

= Protocols NVMe subsystem-maps endpoint overview

== Overview

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

The NVMe subsystem map REST API allows you to create, delete and discover NVMe subsystem maps.

== Examples

=== Creating an NVMe subsystem map

The API:

POST /api/protocols/nvme/subsystem-maps

The call:

```
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystem-maps' -H
'accept: application/hal+json' -d '{ "svm": { "name": "svm1" },
"subsystem": { "name": "subsystem1" }, "namespace": { "name":
"/vol/vol1/namespace1" } }'
```

'''

=== Retrieving all of the NVMe subsystem maps

The API:

GET /api/protocols/nvme/subsystem-maps

The call:

curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-maps' -H
'accept: application/hal+json'

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "0e91b214-fe40-11e8-91a0-005056a79967",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/0e91b214-fe40-11e8-91a0-005056a79967"
          }
        }
      },
      "subsystem": {
        "uuid": "580a6b1e-fe43-11e8-91a0-005056a79967",
        "name": "subsystem1",
        "_links": {
          "self": {
            "href": "/api/protocols/nvme/subsystems/580a6b1e-fe43-11e8-91a0-005056a79967"
          }
        }
      },
      "namespace": {
        "uuid": "3ccdedc6-2519-4206-bc1f-b0f4adab6f89",
        "name": "/vol/vol1/namespacel",
        "_links": {
          "self": {
            "href": "/api/storage/namespaces/3ccdedc6-2519-4206-bc1f-b0f4adab6f89"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/protocols/nvme/subsystem-maps/580a6b1e-fe43-11e8-
```

```
91a0-005056a79967/3ccdedc6-2519-4206-bc1f-b0f4adab6f89"
```

```
    }  
  }  
}  
],  
"num_records": 1,  
"_links": {  
  "self": {  
    "href": "/api/protocols/nvme/subsystem-maps"  
  }  
}  
}  
}
```

'''

=== Retrieving a specific NVMe subsystem map

The NVMe subsystem map is identified by the UUID of the NVMe subsystem followed by the UUID of the NVMe namespace.

The API:

```
GET /api/protocols/nvme/subsystem-maps/{subsystem.uuid}/{namespace.uuid}
```

The call:

```
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystem-maps/580a6b1e-  
fe43-11e8-91a0-005056a79967/3ccdedc6-2519-4206-bc1f-b0f4adab6f89' -H  
'accept: application/hal+json'
```

The response:

```
{  
  "svm": {  
    "uuid": "0e91b214-fe40-11e8-91a0-005056a79967",  
    "name": "svm1",  
    "_links": {  
      "self": {  
        "href": "/api/svm/svms/0e91b214-fe40-11e8-91a0-005056a79967"  
      }  
    }  
  },  
  "subsystem": {  
    "uuid": "580a6b1e-fe43-11e8-91a0-005056a79967",  
    "name": "subsystem1",  
    "_links": {
```

```

    "self": {
      "href": "/api/protocols/nvme/subsystems/580a6b1e-fe43-11e8-91a0-005056a79967"
    }
  },
  "namespace": {
    "uuid": "3ccdedc6-2519-4206-bc1f-b0f4adab6f89",
    "name": "/vol/vol1/namespace1",
    "node": {
      "name": "node1",
      "uuid": "012b4508-67d6-4788-8c2d-801f254ce976",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/012b4508-67d6-4788-8c2d-801f254ce976"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/namespaces/3ccdedc6-2519-4206-bc1f-b0f4adab6f89"
      }
    }
  },
  "nsid": "00000001h",
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystem-maps/580a6b1e-fe43-11e8-91a0-005056a79967/3ccdedc6-2519-4206-bc1f-b0f4adab6f89"
    }
  }
}

```

'''

=== Deleting an NVMe subsystem map

The API:

```
DELETE /api/protocols/nvme/subsystem-
maps/{subsystem.uuid}/{namespace.uuid}
```

The call:

```
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystem-  
maps/580a6b1e-fe43-11e8-91a0-005056a79967/3ccdedc6-2519-4206-bc1f-  
b0f4adab6f89' -H 'accept: application/hal+json'  
----
```

```
[[IDe1618236ce0145a12ad65937d3e6a722]]  
= Retrieve NVMe subsystem maps
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/protocols/nvme/subsystem-maps`#
```

Introduced In: 9.6

Retrieves NVMe subsystem maps.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `anagrpid`

== Related ONTAP commands

* `vserver nvme subsystem map show`

== Learn more

* [xref:{relative_path}protocols_nvme_subsystem-maps_endpoint_overview.html\[DOC /protocols/nvme/subsystem-maps\]](#)

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In
```

```
|Required
|Description

|namespace.uuid
|string
|query
|False
a|Filter by namespace.uuid


|namespace.name
|string
|query
|False
a|Filter by namespace.name


|namespace.node.name
|string
|query
|False
a|Filter by namespace.node.name


|namespace.node.uuid
|string
|query
|False
a|Filter by namespace.node.uuid


|subsystem.name
|string
|query
|False
a|Filter by subsystem.name


|subsystem.uuid
|string
|query
|False
a|Filter by subsystem.uuid


|svm.uuid
|string
```

```
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|anagrpId
|string
|query
|False
a|Filter by anagrpId
```

```
|nsid
|string
|query
|False
a|Filter by nsid
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

|return_timeout
|integer
|query
|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]

a|

|num_records
|integer

a|Number of records.

|records


```
|array[link:#nvme_subsystem_map[nvme_subsystem_map]]
```

```
a|
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anagrpId": "00103050h",
      "namespace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "/vol/vol1/namespacel",
        "node": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h",

```

```

    "subsystem": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{

```

```

"error": {
  "arguments": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

|===

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

```

```
|uuid
|string
a|
```

```
|===
```

```
[#namespace]
[.api-collapsible-fifth-title]
namespace
```

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The fully qualified path name of the NVMe namespace composed from the
volume name, qtree name, and file name of the NVMe namespace. Valid in
POST.
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|The unique identifier of the NVMe namespace. Valid in POST.
```

```
|===
```

```
[#subsystem]
[.api-collapsible-fifth-title]
```

subsystem

The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the NVMe subsystem.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_subsystem_map]
[.api-collapsible-fifth-title]
nvme_subsystem_map
```

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|anagrpId
|string
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe
namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|namespace
|link:#namespace[namespace]
```

a|The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

|nsid

|string

a|The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

|subsystem

|link:#subsystem[subsystem]

a|The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

|svm

|link:#svm[svm]

a|

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===


```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID88f06fa47ca85701d650435b9859c215]]
= Create an NVMe subsystem map

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/nvme/subsystem-maps`#

*Introduced In:* 9.6

```

Creates an NVMe subsystem map.

== Required properties

- * `svm.uuid` or `svm.name` - Existing SVM in which to create the NVMe subsystem map.
- * `namespace.uuid` or `namespace.name` - Existing NVMe namespace to map to the specified NVMe subsystem.
- * `subsystem.uuid` or `subsystem.name` - Existing NVMe subsystem to map to the specified NVMe namespace.

== Related ONTAP commands

- * `vserver nvme subsystem map create`

== Learn more

- * [xref:{relative_path}protocols_nvme_subsystem-maps_endpoint_overview.html](#)[DOC /protocols/nvme/subsystem-maps]

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

return_records
boolean
query
False

a|The default is false. If set to true, the records are returned.

- * Default value:

|===

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|anagrpid
```

```
|string
```

```
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|namespace
```

```
|link:#namespace[namespace]
```

```
a|The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.
```

```
|nsid
```

```
|string
```

```
a|The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.
```

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|subsystem
```

```
|link:#subsystem[subsystem]
```

```
a|The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "anagrpId": "00103050h",
  "namespace": {
    "name": "/vol/vol1/namespace1",
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "nsid": "00000001h",
  "subsystem": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====
```

== Response

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_subsystem_map[nvme_subsystem_map]]
```

```

a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "anagrpId": "00103050h",
      "namespace": {
        "name": "/vol/vol1/namespace1",
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "000000001h",
      "subsystem": {
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 72089790
| The supplied NVMe namespace is already mapped to the supplied NVMe

```

```
subsystem.  
  
| 72089793  
| An NVMe namespace in a Snapshot copy cannot be mapped.  
  
| 72089799  
| The NVMe namespace is the destination of an ongoing restore operation  
and is inaccessible for I/O and management.  
  
| 72089902  
| A node does not have an NVMe interface configured.  
  
| 72089903  
| Multiple nodes do not have an NVMe interface configured.  
  
| 72089904  
| The aggregate must be given back to its home node prior to mapping the  
NVMe namespace it contains.  
  
| 72090001  
| The NVMe subsystem specified by `subsystem.uuid` was not found.  
  
| 72090005  
| The specified `namespace.uuid` and `namespace.name` refer to different  
NVMe namespaces.  
  
| 72090006  
| The NVMe namespace specified by `namespace.uuid` was not found.  
  
| 72090007  
| The NVMe namespace specified by `namespace.name` was not found.  
  
| 72090020  
| The specified `subsystem.uuid` and `subsystem.name` refer to different  
NVMe subsystems.  
  
| 72090021  
| The NVMe subsystem specified by `subsystem.name` was not found.  
|==
```

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#namespace]
[.api-collapsible-fifth-title]
namespace

```

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST

by supplying either the UUID, name, or both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|The fully qualified path name of the NVMe namespace composed from the volume name, qtree name, and file name of the NVMe namespace. Valid in POST.

```
|node
```

```
|link:#node[node]
```

a|

```
|uuid
```

```
|string
```

a|The unique identifier of the NVMe namespace. Valid in POST.

```
|===
```

```
[#subsystem]
```

```
[.api-collapsible-fifth-title]
```

subsystem

The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|The name of the NVMe subsystem.


```
|uuid
|string
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_subsystem_map]
[.api-collapsible-fifth-title]
nvme_subsystem_map
```

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

|anagrpid
|string
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|namespace
|link:#namespace[namespace]
a|The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

|nsid
|string
a|The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[IDa0262c67f0743e29c0e01c9606e59d65]]

= Delete an NVMe subsystem map

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/protocols/nvme/subsystem-maps/{subsystem.uuid}/{namespace.uuid}`#

Introduced In: 9.6

Deletes an NVMe subsystem map.

== Related ONTAP commands

* `vserver nvme subsystem map delete`

== Learn more

* xref:{relative_path}protocols_nvme_subsystem-maps_endpoint_overview.html[DOC /protocols/nvme/subsystem-maps]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|subsystem.uuid

|string

|path

|True

a|The unique identifier of the NVMe subsystem.

```

|namespace.uuid
|string
|path
|True
a|The unique identifier of the NVMe namespace.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_subsystem_map[nvme_subsystem_map]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {

```

```

    "href": "/api/resourcelink"
  }
},
"records": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "anagrpId": "00103050h",
    "namespace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/vol1/namespacel",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "nsid": "000000001h",
    "subsystem": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",

```

```

        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 72090019
| The specified NVMe namespace is not mapped to the specified NVMe
subsystem.
|===

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ]
  }
}

```

```

    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next
|link:#href[href]
a|

```



```

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

```

```
[#namespace]
[.api-collapsible-fifth-title]
namespace
```

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The fully qualified path name of the NVMe namespace composed from the
volume name, qtree name, and file name of the NVMe namespace. Valid in
POST.
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|The unique identifier of the NVMe namespace. Valid in POST.
```

```
|===
```

```
[#subsystem]
[.api-collapsible-fifth-title]
subsystem
```

The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the NVMe subsystem.

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

```
|===
```

```
[#nvme_subsystem_map]  
[.api-collapsible-fifth-title]  
nvme_subsystem_map
```

An NVMe subsystem map is an association of an NVMe namespace with an NVMe subsystem. When an NVMe namespace is mapped to an NVMe subsystem, the NVMe subsystem's hosts are granted access to the NVMe namespace. The relationship between an NVMe subsystem and an NVMe namespace is one subsystem to many namespaces.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
 |_links  
|link:#_links[_links]  
a|
```

```
|anagrpid  
|string  
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe  
namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|namespace  
|link:#namespace[namespace]  
a|The NVMe namespace to which the NVMe subsystem is mapped. Required in  
POST by supplying either the UUID, name, or both.
```

```
|nsid  
|string
```

a|The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

|subsystem

|link:#subsystem[subsystem]

a|The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

|svm

|link:#svm[svm]

a|

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDff09d1d38e7e4d008f1e735ac2532326]]
= Retrieve an NVMe subsystem map

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nvme/subsystem-maps/{subsystem.uuid}/{namespace.uuid}`#

*Introduced In:* 9.6

Retrieves an NVMe subsystem map.

== Expensive properties

There is an added cost to retrieving values for these properties. They are

```

not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `anagrpid`

== Related ONTAP commands

* `vserver nvme subsystem map show`

== Learn more

* [xref:{relative_path}protocols_nvme_subsystem-maps_endpoint_overview.html\[DOC /protocols/nvme/subsystem-maps\]](#)

== Parameters

[cols=5*,options=header]
|==

|Name
|Type
|In
|Required
|Description

|subsystem.uuid
|string
|path
|True
a|The unique identifier of the NVMe subsystem.

|namespace.uuid
|string
|path
|True
a|The unique identifier of the NVMe namespace.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|anagrpid
```

```
|string
```

```
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|namespace
```

```
|link:#namespace[namespace]
```

```
a|The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.
```

```
|nsid
```

```
|string
```

```
a|The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.
```

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".


```
|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe namespace is mapped. Required in
POST by supplying either `subsystem.uuid`, `subsystem.name` or both.
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anagrpid": "00103050h",
  "namespace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "/vol/vol1/namespace1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "nsid": "00000001h",
  "subsystem": {
    "_links": {
```

```

    "self": {
      "href": "/api/resourcelink"
    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 72090019
| The specified NVMe namespace is not mapped to the specified NVMe
subsystem.
|===

```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```
|===
```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]

```

```

[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#namespace]
[.api-collapsible-fifth-title]
namespace

```

The NVMe namespace to which the NVMe subsystem is mapped. Required in POST by supplying either the UUID, name, or both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The fully qualified path name of the NVMe namespace composed from the  
volume name, qtree name, and file name of the NVMe namespace. Valid in  
POST.
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the NVMe namespace. Valid in POST.
```

```
|===
```

```
[#subsystem]
```

```
[.api-collapsible-fifth-title]
```

```
subsystem
```

The NVMe subsystem to which the NVMe namespace is mapped. Required in POST by supplying either `subsystem.uuid`, `subsystem.name` or both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```

a|

|name
|string
a|The name of the NVMe subsystem.

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```
|===
```

```
//end collapsible .Definitions block  
=====
```

```
:leveloffset: -1
```

```
= Manage NVMe subsystems
```

```
:leveloffset: +1
```

```
[[ID02f06c7c6d4ce42f950063c6106fb310]]  
= Protocols NVMe subsystems endpoint overview
```

```
== Overview
```

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

The NVMe subsystem REST API allows you to create, update, delete, and discover NVMe subsystems as well as add and remove NVMe hosts that can access the subsystem and associated namespaces.

```
== Examples
```

```
=== Creating an NVMe subsystem
```

```
----
```

```
# The API:  
POST /api/protocols/nvme/subsystems
```

```
# The call:  
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystems' -H 'accept:  
application/json' -d '{ "svm": { "name": "svm1" }, "name": "subsystem1",  
"os_type": "linux" }'  
----
```

```
'''
```



```
=== Creating an NVMe subsystem with multiple NVMe subsystem hosts
```

```
----
```

```
# The API:
```

```
POST /api/protocols/nvme/subsystems
```

```
# The call:
```

```
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystems' -H 'accept: application/json' -d '{ "svm": { "name": "svm1" }, "name": "subsystem2", "os_type": "vmware", "hosts": [ { "nqn": "nqn.1992-01.example.com:host1" }, { "nqn": "nqn.1992-01.example.com:host2" } ] }'
```

```
----
```

```
'''
```

```
=== Retrieving all NVMe subsystems
```

```
----
```

```
# The API:
```

```
GET /api/protocols/nvme/subsystems
```

```
# The call:
```

```
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems' -H 'accept: application/json'
```

```
# The response:
```

```
{
  "records": [
    {
      "svm": {
        "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
        "name": "svm1",
      },
      "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
      "name": "subsystem1",
    },
    {
      "svm": {
        "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
        "name": "svm1",
      },
      "uuid": "bcde901a-a379-4a91-9ea6-1b728ed6696f",
      "name": "subsystem2",
    }
  ]
}
```

```

    }
],
"num_records": 2,
}
----

'''

=== Retrieving all NVMe subsystems with OS type _linux_

Note that the `os_type` query parameter is used to perform the query.

----

# The API:
GET /api/protocols/nvme/subsystems

# The call:
curl -X GET 'https://<mgmt-
ip>/api/protocols/nvme/subsystems?os_type=linux' -H 'accept:
application/json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
      "name": "svm1",
    },
    "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
    "name": "subsystem1",
    "os_type": "linux",
  }
],
"num_records": 1,
}
----

'''

=== Retrieving a specific NVMe subsystem

----

# The API:
GET /api/protocols/nvme/subsystems/{uuid}

```

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f' -H 'accept: application/json'

# The response:
{
  "svm": {
    "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
    "name": "svml",
  },
  "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
  "name": "subsystem1",
  "os_type": "linux",
  "target_nqn": "nqn.1992-
08.com.netapp:sn.d04594ef915b4c73b642169e72e4c0b1:subsystem.subsystem1",
  "serial_number": "wtJNKNKD-uPLAAAAAAD",
  "io_queue": {
    "default": {
      "count": 4,
      "depth": 32
    }
  }
}
}
-----

'''
```

=== Retrieving the NVMe namespaces mapped to a specific NVMe subsystem

Note that the `fields` query parameter is used to specify the desired properties.

```
-----

# The API:
GET /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f?fields=subsystem_maps' -H 'accept:
application/json'

# The response:
{
  "svm": {
```

```

    "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
    "name": "svml",
  },
  "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
  "name": "subsystem1",
  "subsystem_maps": [
    {
      "anagrp_id": "00000001h",
      "namespace": {
        "uuid": "eeaaca23-128d-4a7d-be4a-dc9106705799",
        "name": "/vol/vol1/namespace1"
      },
      "nsid": "00000001h"
    },
    {
      "anagrp_id": "00000002h",
      "namespace": {
        "uuid": "feaaca23-83a0-4a7d-beda-dc9106705799",
        "name": "/vol/vol1/namespace2"
      },
      "nsid": "00000002h"
    }
  ]
}
----

'''

=== Adding a comment about an NVMe subsystem

----

# The API:
PATCH /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f' -H 'accept: application/json' -d '{
"comment": "A brief comment about the subsystem" }'
----

'''

=== Deleting an NVMe subsystem

----

```

```
# The API:
DELETE /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f' -H 'accept: application/json'
----
```

=== Deleting an NVMe subsystem with mapped NVMe namespaces

Normally, deleting an NVMe subsystem that has mapped NVMe namespaces is not allowed. The deletion can be forced using the `allow_delete_while_mapped` query parameter.

```
-----

# The API:
DELETE /api/protocols/nvme/subsystems/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f?allow_delete_while_mapped=true' -H 'accept:
application/json'
-----

=== Delete an NVMe subsystem with NVMe subsystem hosts


Normally, deleting an NVMe subsystem with NVMe subsystem hosts is disallowed. The deletion can be forced using the `allow_delete_with_hosts` query parameter.



```

The API:
DELETE /api/protocols/nvme/subsystems/{uuid}

The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f?allow_delete_with_hosts=true' -H 'accept:
application/json'

'''

== An NVMe Subsystem Host
```



An NVMe subsystem host is a network host provisioned to an NVMe subsystem


```

to access namespaces mapped to that subsystem.

== Examples

=== Adding an NVMe subsystem host to an NVMe subsystem

The API:

POST /protocols/nvme/subsystems/{subsystem.uuid}/hosts

The call:

```
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-  
a379-4a91-9ea6-1b728ed6696f/hosts' -H 'accept: application/json' -d '{  
  "nqn": "nqn.1992-01.com.example:subsys1.host1" }'
```

'''

=== Adding multiple NVMe subsystem hosts to an NVMe subsystem

The API:

POST /protocols/nvme/subsystems/{subsystem.uuid}/hosts

The call:

```
curl -X POST 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-  
a379-4a91-9ea6-1b728ed6696f/hosts' -H 'accept: application/json' -d '{  
  "records": [ { "nqn": "nqn.1992-01.com.example:subsys1.host2" }, { "nqn":  
    "nqn.1992-01.com.example:subsys1.host3" } ] }'
```

'''

=== Retrieving all NVMe subsystem hosts for an NVMe subsystem

The API:

GET /protocols/nvme/subsystems/{subsystem.uuid}/hosts

The call:

```
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-  
a379-4a91-9ea6-1b728ed6696f/hosts' -H 'accept: application/json'
```

The response:

```
{
"records": [
  {
    "nqn": "nqn.1992-01.com.example:subsys1.host1",
  },
  {
    "nqn": "nqn.1992-01.com.example:subsys1.host2",
  },
  {
    "nqn": "nqn.1992-01.com.example:subsys1.host3",
  }
],
"num_records": 3,
}
----
```

'''

=== Retrieving a specific NVMe subsystem host for an NVMe subsystem

The API:

GET /protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}

The call:

```
curl -X GET 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-
a379-4a91-9ea6-1b728ed6696f/hosts/nqn.1992-01.com.example:subsys1.host1'
-H 'accept: application/json'
```

The response:

```
{
"subsystem": {
  "uuid": "acde901a-a379-4a91-9ea6-1b728ed6696f",
},
"nqn": "nqn.1992-01.com.example:subsys1.host1",
"io_queue": {
  "count": 4,
  "depth": 32
},
}
----
```

'''

=== Deleting an NVMe subsystem host from an NVMe subsystem

The API:

DELETE /protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}

The call:

```
curl -X DELETE 'https://<mgmt-ip>/api/protocols/nvme/subsystems/acde901a-  
a379-4a91-9ea6-1b728ed6696f/hosts/nqn.1992-01.com.example:subsys1.host1'  
-H 'accept: application/json'
```

[[ID69b7bbfcde3fbf38460e7a270cfb1f0a]]

= Retrieve NVMe subsystems

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nvme/subsystems`#

Introduced In: 9.6

Retrieves NVMe subsystems.

== Related ONTAP commands

* `vserver nvme subsystem host show`

* `vserver nvme subsystem map show`

* `vserver nvme subsystem show`

== Learn more

* xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html[DOC
/protocols/nvme/subsystems]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description


```
|os_type
|string
|query
|False
a|Filter by os_type
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|hosts.nqn
|string
|query
|False
a|Filter by hosts.nqn
```

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|target_nqn
|string
|query
|False
a|Filter by target_nqn
```

```
|subsystem_maps.anagrpid
|string
|query
|False
```

```
a|Filter by subsystem_maps.anagrpid

|subsystem_maps.nsid
|string
|query
|False
a|Filter by subsystem_maps.nsid

|subsystem_maps.namespace.name
|string
|query
|False
a|Filter by subsystem_maps.namespace.name

|subsystem_maps.namespace.uuid
|string
|query
|False
a|Filter by subsystem_maps.namespace.uuid

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|serial_number
|string
|query
|False
a|Filter by serial_number

|io_queue.default.depth
|integer
```

```
|query
|False
a|Filter by io_queue.default.depth
```

```
|io_queue.default.count
|integer
|query
|False
a|Filter by io_queue.default.count
```

```
|delete_on_unmap
|boolean
|query
|False
a|Filter by delete_on_unmap
```

* Introduced in: 9.7

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
```

```

|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_subsystem[nvme_subsystem]]
a|

|===

```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "hosts": [
        {
          "nqn": "nqn.1992-01.example.com:string"
        }
      ],
      "io_queue": {
        "default": {
          "count": "4",
          "depth": "16"
        }
      },
      "name": "subsystem1",
      "os_type": "string",
      "serial_number": "wCVsgFMiuMhVAAAAAAB",
      "subsystem_maps": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "anagrpid": "00103050h",
          "namespace": {
            "_links": {
```

```

        "self": {
            "href": "/api/resourcelink"
        },
        "name": "/vol/vol1/namespace1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "nsid": "00000001h"
}
],
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"target_nqn": "nqn.1992-01.example.com:string",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]

```

```

=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#hosts]
[.api-collapsible-fifth-title]
hosts

[cols=3*,options=header]
|===
|Name
|Type
|Description

|nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.

```


|===

```
[#default]
[.api-collapsible-fifth-title]
default
```

The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of host I/O queue pairs.
```

```
|depth
|integer
a|The host I/O queue depth.
```

|===

```
[#io_queue]
[.api-collapsible-fifth-title]
io_queue
```

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|default
|link:#default[default]
```

a|The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

|===

[#namespace]
[.api-collapsible-fifth-title]
namespace

An NVMe namespace mapped to the NVMe subsystem.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]

a|

|name
|string

a|The name of the NVMe namespace.

|uuid
|string

a|The unique identifier of the NVMe namespace.

|===

[#subsystem_maps]
[.api-collapsible-fifth-title]
subsystem_maps

An NVMe namespace mapped to the NVMe subsystem.

[cols=3*,options=header]

|===

|Name

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|anagrpId
|string
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe
namespace.

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed
by a lower case "h".

|namespace
|link:#namespace[namespace]
a|An NVMe namespace mapped to the NVMe subsystem.

|nsid
|string
a|The NVMe namespace identifier. This is an identifier used by an NVMe
controller to provide access to the NVMe namespace.

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-
filled) followed by a lower case "h".

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_subsystem]
[.api-collapsible-fifth-title]
nvme_subsystem
```

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
 |_links
|link:#_links[_links]
a|
```

```
|comment
|string
a|A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
```

```
|delete_on_unmap
|boolean
a|An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. This property defaults to _false_ when the subsystem is created.
```

```
|hosts
|array[link:#hosts[hosts]]
a|The NVMe hosts configured for access to the NVMe subsystem. Optional in
```

POST.

|io_queue

|link:#io_queue[io_queue]

a|The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

|name

|string

a|The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.

|os_type

|string

a|The host operating system of the NVMe subsystem's hosts. Required in POST.

|serial_number

|string

a|The serial number of the NVMe subsystem.

|subsystem_maps

|array[link:#subsystem_maps[subsystem_maps]]

a|The NVMe namespaces mapped to the NVMe subsystem.

There is an added cost to retrieving property values for `subsystem_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|svm

|link:#svm[svm]

a|

|target_nqn

|string

a|The NVMe qualified name (NQN) used to identify the NVMe storage target.

|uuid

```
|string
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[IDc8e8e57bb6cf8e222d83fcb353184be4]]

= Create an NVMe subsystem

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/protocols/nvme/subsystems`#

Introduced In: 9.6

Creates an NVMe subsystem.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the NVMe subsystem.

* `name` - Name for NVMe subsystem. Once created, an NVMe subsystem cannot be renamed.

* `os_type` - Operating system of the NVMe subsystem's hosts.

== Related ONTAP commands

* `vserver nvme subsystem create`

== Learn more

* xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html[DOC /protocols/nvme/subsystems]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|comment

|string

a|A configurable comment for the NVMe subsystem. Optional in POST and PATCH.

|delete_on_unmap

|boolean

a|An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. This property defaults to `_false_` when the subsystem is created.

|hosts

|array[link:#hosts[hosts]]

a|The NVMe hosts configured for access to the NVMe subsystem. Optional in POST.

|io_queue

|link:#io_queue[io_queue]

a|The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

|name

|string

a|The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.

|os_type

|string

a|The host operating system of the NVMe subsystem's hosts. Required in POST.

|serial_number

|string

a|The serial number of the NVMe subsystem.

|subsystem_maps

|array[link:#subsystem_maps[subsystem_maps]]

a|The NVMe namespaces mapped to the NVMe subsystem.

There is an added cost to retrieving property values for `subsystem_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|svm

|link:#svm[svm]

a|

|target_nqn

|string

a|The NVMe qualified name (NQN) used to identify the NVMe storage target.

```

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "comment": "string",
  "hosts": [
    {
      "nqn": "nqn.1992-01.example.com:string"
    }
  ],
  "io_queue": {
    "default": {
      "count": "4",
      "depth": "16"
    }
  },
  "name": "subsystem1",
  "os_type": "string",
  "serial_number": "wCVsgFMiuMhVAAAAAAB",
  "subsystem_maps": [
    {
      "anagrpid": "00103050h",
      "namespace": {
        "name": "/vol/vol1/namespace1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target_nqn": "nqn.1992-01.example.com:string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

```

```
====
```

== Response

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_subsystem[nvme_subsystem]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "comment": "string",
      "hosts": [
        {
          "nqn": "nqn.1992-01.example.com:string"
        }
      ],
      "io_queue": {
        "default": {
          "count": "4",
          "depth": "16"
        }
      },
      "name": "subsystem1",
      "os_type": "string",
      "serial_number": "wCVsgFMiuMhVAAAAAAB",
```

```

    "subsystem_maps": [
      {
        "anagrpid": "00103050h",
        "namespace": {
          "name": "/vol/vol1/namespace1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "nsid": "00000001h"
      }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "target_nqn": "nqn.1992-01.example.com:string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 2621462
| The supplied SVM does not exist.
|
| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.
|
| 2621707
| The svm.uuid or svm.name must be provided.
|
| 72089709
| The NVMe subsystem name contains an invalid character.
|
| 72089771
| The NQN is invalid. A non-empty qualifier is required after the prefix.
An example of a valid NQN is _nqn.1992-01.com.example:string_.

```

```

| 72089772
| The NQN is invalid. Add the prefix 'nqn'. An example of a valid NQN is
_nqn.1992-01.com.example:string_.

| 72089773
| The NQN is invalid. The date field must be formatted _yyyy-mm_. An
example of a valid NQN is _nqn.1992-01.com.example:string_.

| 72090025
| The NVMe subsystem already exists for the SVM.

| 72090029
| The NVMe service does not exist.

| 72090030
| A partial success occurred while adding multiple NVMe subsystem hosts to
an NVMe subsystem.

| 72090035
| Passing NVMe subsystem host NQNs on NVMe subsystem POST requires an
effective cluster version of 9.7 or later.

| 72090036
| The `hosts.nqn` NVMe subsystem property must contain unique values.
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

```

|===

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#hosts]
[.api-collapsible-fifth-title]
hosts
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|nqn
```

```
|string
```

a|The NVMe qualified name (NQN) used to identify the NVMe storage target.

|===

```
[#default]
[.api-collapsible-fifth-title]
default
```

The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|count
```

```
|integer
```

a|The number of host I/O queue pairs.

```
|depth
```

```
|integer
```

a|The host I/O queue depth.

|===

```
[#io_queue]
[.api-collapsible-fifth-title]
io_queue
```

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|default
```

```
|link:#default[default]
```

a|The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

|===

```
[#namespace]
[.api-collapsible-fifth-title]
namespace
```

An NVMe namespace mapped to the NVMe subsystem.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|name
```

```
|string
```

a|The name of the NVMe namespace.

```
|uuid
```

```
|string
a|The unique identifier of the NVMe namespace.
```

```
|===
```

```
[#subsystem_maps]
[.api-collapsible-fifth-title]
subsystem_maps
```

An NVMe namespace mapped to the NVMe subsystem.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|anagrpid
```

```
|string
```

```
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe
namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|namespace
```

```
|link:#namespace[namespace]
```

```
a|An NVMe namespace mapped to the NVMe subsystem.
```

```
|nsid
```

```
|string
```

```
a|The NVMe namespace identifier. This is an identifier used by an NVMe
controller to provide access to the NVMe namespace.
```

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|===
```

```
[#svm]
```



```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_subsystem]
```

```
[.api-collapsible-fifth-title]
```

```
nvme_subsystem
```

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|comment
```

```
|string
```

```
a|A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
```

```
|delete_on_unmap
```

```
|boolean
```

```
a|An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. This property defaults to _false_ when the subsystem is created.
```

|hosts
|array[link:#hosts[hosts]]
a|The NVMe hosts configured for access to the NVMe subsystem. Optional in POST.

|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

|name
|string
a|The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.

|os_type
|string
a|The host operating system of the NVMe subsystem's hosts. Required in POST.

|serial_number
|string
a|The serial number of the NVMe subsystem.

|subsystem_maps
|array[link:#subsystem_maps[subsystem_maps]]
a|The NVMe namespaces mapped to the NVMe subsystem.

There is an added cost to retrieving property values for `subsystem_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|svm
|link:#svm[svm]
a|

|target_nqn

```

|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]

```

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[IDa9a1e00517f1d53ca4847916e370bf4b]]

= Retrieve NVMe subsystem hosts

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nvme/subsystems/{subsystem.uuid}/hosts`#

Introduced In: 9.6

Retrieves the NVMe subsystem hosts of an NVMe subsystem.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `subsystem_maps`.*+`

== Related ONTAP commands

```
* `vserver nvme subsystem map show`
```

```
* `vserver nvme subsystem show`
```

== Learn more

```
* xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html[DOC  
/protocols/nvme/subsystems]
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|subsystem.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|The unique identifier of the NVMe subsystem.
```

```
|fields
```

```
|array[string]
```

```
|query
```

```
|False
```

```
a|Specify the fields to return.
```

```
|max_records
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Limit the number of records returned.
```

```
|return_records
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|The default is true for GET calls. When set to false, only the number
```

of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

```
|records
|array[link:#nvme_subsystem_host[nvme_subsystem_host]]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

```
=====
```

[source,json,subs==+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "io_queue": {
        "count": "4",
        "depth": "32"
      },
      "nqn": "nqn.1992-01.example.com:string",
      "subsystem": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
```

```
=====
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```



```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]

```

a|

|===

[#io_queue]
[.api-collapsible-fifth-title]
io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem.

|depth
|integer
a|The I/O queue depth. The default value is inherited from the owning NVMe subsystem.

|===

[#subsystem]
[.api-collapsible-fifth-title]
subsystem

The NVMe subsystem to which the NVMe host has been provisioned.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===

[#records]
[.api-collapsible-fifth-title]
records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for
execution by the NVMe controller.

|nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.
Not allowed in POST when the `records` property is used.

|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe host has been provisioned.

|===

```

```
[#nvme_subsystem_host]
[.api-collapsible-fifth-title]
nvme_subsystem_host
```

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for
execution by the NVMe controller.
```

```
|nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.
Not allowed in POST when the `records` property is used.
```

```
|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe host has been provisioned.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

```
//end collapsible .Definitions block
====
```

```
[[IDe1ble4e1a4066135d916060cdba3f23b]]
= Add NVMe subsystem hosts
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/nvme/subsystems/{subsystem.uuid}/hosts`#
```

Introduced In: 9.6

Adds NVMe subsystem host(s) to an NVMe subsystem.

== Required properties

* `nqn` or `records.nqn` - NVMe host(s) NQN(s) to add to the NVMe subsystem.

== Related ONTAP commands

* `vserver nvme subsystem host add`

== Learn more

* xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html[DOC /protocols/nvme/subsystems]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|subsystem.uuid
|string
|path
|True
```

```
a|The unique identifier of the NVMe subsystem.
```

```

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for
execution by the NVMe controller.

|nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.
Not allowed in POST when the `records` property is used.

|records
|array[link:#records[records]]
a|An array of NVMe hosts specified to add multiple NVMe hosts to an NVMe
subsystem in a single API call. Valid in POST only.

|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe host has been provisioned.

|===

```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "io_queue": {
    "count": "4",
    "depth": "32"
  },
  "nqn": "nqn.1992-01.example.com:string",
  "records": [
    {
      "io_queue": {
        "count": "4",
        "depth": "32"
      },
      "nqn": "nqn.1992-01.example.com:string",
      "subsystem": {
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "subsystem": {
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
```



```
|array[link:#nvme_subsystem_host[nvme_subsystem_host]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "io_queue": {
        "count": "4",
        "depth": "32"
      },
      "nqn": "nqn.1992-01.example.com:string",
      "records": [
        {
          "io_queue": {
            "count": "4",
            "depth": "32"
          },
          "nqn": "nqn.1992-01.example.com:string",
          "subsystem": {
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        }
      ],
      "subsystem": {
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
====

== Error
```

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
72089705	The NVMe subsystem host already exists for the NVMe subsystem.
72089771	The NQN is invalid. A non-empty qualifier is required after the prefix. An example of a valid NQN is <code>_nqn.1992-01.com.example:string_</code> .
72089772	The NQN is invalid. Add the prefix <code>_nqn_</code> . An example of a valid NQN is <code>_nqn.1992-01.com.example:string_</code> .
72089773	The NQN is invalid. The date field must be formatted <code>_yyyy-mm_</code> . An example of a valid NQN is <code>_nqn.1992-01.com.example:string_</code> .
72090002	The POST request of hosts to an NVMe subsystem can only contain an 'nqn' property or 'records' property, but not both.
72090003	The elements in the records array for a POST of hosts to an NVMe subsystem must contain only the nqn property.

===

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|
```

|===

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#io_queue]
[.api-collapsible-fifth-title]
io_queue
```

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|count
|integer
```

a|The number of I/O queue pairs. The default value is inherited from the owning NVMe subsystem.

```
|depth
|integer
```

a|The I/O queue depth. The default value is inherited from the owning NVMe subsystem.

|===

```
[#subsystem]
[.api-collapsible-fifth-title]
subsystem
```

The NVMe subsystem to which the NVMe host has been provisioned.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
```

```

|Description

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===

[#records]
[.api-collapsible-fifth-title]
records

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for
execution by the NVMe controller.

|nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.
Not allowed in POST when the `records` property is used.

|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe host has been provisioned.

|===

[#nvme_subsystem_host]
[.api-collapsible-fifth-title]
nvme_subsystem_host

```

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|io_queue
```

```
|link:#io_queue[io_queue]
```

a|The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

```
|nqn
```

```
|string
```

a|The NVMe qualified name (NQN) used to identify the NVMe storage target. Not allowed in POST when the `records` property is used.

```
|records
```

```
|array[link:#records[records]]
```

a|An array of NVMe hosts specified to add multiple NVMe hosts to an NVMe subsystem in a single API call. Valid in POST only.

```
|subsystem
```

```
|link:#subsystem[subsystem]
```

a|The NVMe subsystem to which the NVMe host has been provisioned.

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

error_arguments

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[ID247e6a8164cfbb0fd70efd7d2436f313]]
```

= Delete an NVMe subsystem host

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}`#
```

Introduced In: 9.6

Deletes an NVMe subsystem host from an NVMe subsystem.

== Related ONTAP commands

* `vserver nvme subsystem host remove`

== Learn more

* xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html[DOC /protocols/nvme/subsystems]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|subsystem.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|The unique identifier of the NVMe subsystem.
```

```
|nqn
```

```
|string
```

```
|path
```

```
|True
```

```
a|The NVMe qualified name (NQN) used to identify the NVMe subsystem host.
```

```
|===
```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 72089771

| The NQN is invalid. A non-empty qualifier is required after the prefix.
An example of a valid NQN is `_nqn.1992-01.com.example:string_`.

| 72089772

| The NQN is invalid. Add the prefix `'_nqn'_`. An example of a valid NQN is
`_nqn.1992-01.com.example:string_`.

| 72089773

| The NQN is invalid. The date field must be formatted `_yyyy-mm_`. An
example of a valid NQN is `_nqn.1992-01.com.example:string_`.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====


```

[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]

```

```

[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID911b12839a6af8b50365a5bce8cb3dcc]]
= Retrieve an NVMe subsystem host

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/nvme/subsystems/{subsystem.uuid}/hosts/{nqn}`#

*Introduced In:* 9.6

Retrieves an NVMe subsystem host of an NVMe subsystem.

```

== Related ONTAP commands

* ``vserver nvme subsystem host show``

== Learn more

* `xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html[DOC /protocols/nvme/subsystems]`

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|subsystem.uuid

|string

|path

|True

a|The unique identifier of the NVMe subsystem.

|nqn

|string

|path

|True

a|The NVMe qualified name (NQN) used to identify the NVMe subsystem host.

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for
execution by the NVMe controller.

|nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.
Not allowed in POST when the `records` property is used.

|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe host has been provisioned.

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "io_queue": {
    "count": "4",
    "depth": "32"
  },
  "nqn": "nqn.1992-01.example.com:string",
  "subsystem": {

```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

```
== Definitions
```

```
[.api-def-first-level]
```

```
.See Definitions
```

```
[%collapsible%closed]
```

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#io_queue]
```

```
[.api-collapsible-fifth-title]
```

```
io_queue
```

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|count
```

```
|integer
```

```
a|The number of I/O queue pairs. The default value is inherited from the  
owning NVMe subsystem.
```

```
|depth
```

```
|integer
```

```
a|The I/O queue depth. The default value is inherited from the owning NVMe  
subsystem.
```

```
|===
```

```
[#subsystem]
```

```
[.api-collapsible-fifth-title]
```

```
subsystem
```

The NVMe subsystem to which the NVMe host has been provisioned.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#records]
[.api-collapsible-fifth-title]
records
```

The NVMe host provisioned to access NVMe namespaces mapped to a subsystem.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for
execution by the NVMe controller.
```

```
|nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.
Not allowed in POST when the `records` property is used.
```

```
|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe host has been provisioned.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```



```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[IDed25c32d48c396b36856101e4d94b14a]]
```

```
= Remove an NVMe subsystem
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/protocols/nvme/subsystems/{uuid}`#
```

```
*Introduced In:* 9.6
```

```
Removes an NVMe subsystem.
```

```
== Related ONTAP commands
```

```
* `vserver nvme subsystem delete`
```

```
== Learn more
```

```
* xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html[DOC  
/protocols/nvme/subsystems]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|The unique identifier of the NVMe subsystem.
```

```
|allow_delete_while_mapped
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Allows for the deletion of a mapped NVMe subsystem.
```

```
|allow_delete_with_hosts
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Allows for the deletion of an NVMe subsystem with NVMe hosts.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes
```

```
|===
```

```
| Error Code | Description
```

```
| 72090023
```

```
| The NVMe subsystem contains one or more mapped namespaces. Use the  
`allow_delete_while_mapped` query parameter to delete an NVMe subsystem  
with mapped NVMe namespaces.
```

```
| 72090024
```

```
| The NVMe subsystem contains one or more NVMe hosts. Use the  
`allow_delete_with_hosts` query parameter to delete an NVMe subsystem with  
NVMe hosts.
```

```
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```

a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID12c3a55ad6de6e9626aa3677d1eb7fd7]]

```

= Retrieve an NVMe subsystem

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/nvme/subsystems/{uuid}`#

Introduced In: 9.6

Retrieves an NVMe subsystem.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `subsystem_maps.++`

== Related ONTAP commands

* `vserver nvme subsystem host show`

* `vserver nvme subsystem map show`

* `vserver nvme subsystem show`

== Learn more

* [xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html\[DOC /protocols/nvme/subsystems\]](#)

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|The unique identifier of the NVMe subsystem.

```

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|A configurable comment for the NVMe subsystem. Optional in POST and
PATCH.

|delete_on_unmap
|boolean
a|An option that causes the subsystem to be deleted when the last
subsystem map associated with it is deleted. This property defaults to
_false_ when the subsystem is created.

|hosts
|array[link:#hosts[hosts]]
a|The NVMe hosts configured for access to the NVMe subsystem. Optional in
POST.

|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for

```

execution by the NVMe controller.

|name
|string
a|The name of the NVMe subsystem. Once created, an NVMe subsystem cannot be renamed. Required in POST.

|os_type
|string
a|The host operating system of the NVMe subsystem's hosts. Required in POST.

|serial_number
|string
a|The serial number of the NVMe subsystem.

|subsystem_maps
|array[link:#subsystem_maps[subsystem_maps]]
a|The NVMe namespaces mapped to the NVMe subsystem.

There is an added cost to retrieving property values for `subsystem_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|svm
|link:#svm[svm]
a|

|target_nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===


```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "hosts": [
    {
      "nqn": "nqn.1992-01.example.com:string"
    }
  ],
  "io_queue": {
    "default": {
      "count": "4",
      "depth": "16"
    }
  },
  "name": "subsystem1",
  "os_type": "string",
  "serial_number": "wCVsgFMiuMhVAAAAAAB",
  "subsystem_maps": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "anagrpid": "00103050h",
      "namespace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "/vol/vol1/namespace1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h"
    }
  ],
}
```

```

"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"target_nqn": "nqn.1992-01.example.com:string",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",

```

```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#hosts]

```

[.api-collapsible-fifth-title]

hosts

[cols=3*,options=header]

|===

|Name

|Type

|Description

|nqn

|string

a|The NVMe qualified name (NQN) used to identify the NVMe storage target.

|===

[#default]

[.api-collapsible-fifth-title]

default

The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|count

|integer

a|The number of host I/O queue pairs.

|depth

|integer

a|The host I/O queue depth.

|===

[#io_queue]

[.api-collapsible-fifth-title]

io_queue

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|default
|link:#default[default]
a|The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.
```

```
|===
```

```
[#namespace]
[.api-collapsible-fifth-title]
namespace
```

An NVMe namespace mapped to the NVMe subsystem.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the NVMe namespace.
```

```
|uuid
|string
a|The unique identifier of the NVMe namespace.
```

|===

```
[#subsystem_maps]
[.api-collapsible-fifth-title]
subsystem_maps
```

An NVMe namespace mapped to the NVMe subsystem.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|anagrpid
|string
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe
namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|namespace
|link:#namespace[namespace]
a|An NVMe namespace mapped to the NVMe subsystem.
```

```
|nsid
|string
a|The NVMe namespace identifier. This is an identifier used by an NVMe
controller to provide access to the NVMe namespace.
```

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

|===

```
[#svm]
```

```

[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID59f59bcec1a68144cac918f6d9cbb3]]
= Update an NVMe subsystem

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/nvme/subsystems/{uuid}`#

*Introduced In:* 9.6

```


Updates an NVMe subsystem.

== Related ONTAP commands

* `vserver nvme subsystem modify`

== Learn more

* [xref:{relative_path}protocols_nvme_subsystems_endpoint_overview.html\[DOC /protocols/nvme/subsystems\]](#)

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|The unique identifier of the NVMe subsystem.

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|comment

|string

a|A configurable comment for the NVMe subsystem. Optional in POST and PATCH.

```

|delete_on_unmap
|boolean
a|An option that causes the subsystem to be deleted when the last
subsystem map associated with it is deleted. This property defaults to
_false_ when the subsystem is created.

|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for
execution by the NVMe controller.

|serial_number
|string
a|The serial number of the NVMe subsystem.

|subsystem_maps
|array[link:#subsystem_maps[subsystem_maps]]
a|The NVMe namespaces mapped to the NVMe subsystem.

There is an added cost to retrieving property values for `subsystem_maps`.
They are not populated for either a collection GET or an instance GET
unless explicitly requested using the `fields` query parameter. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

|svm
|link:#svm[svm]
a|

|target_nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===

```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "io_queue": {
    "default": {
      "count": "4",
      "depth": "16"
    }
  },
  "serial_number": "wCVsgFMiuMhVAAAAAAB",
  "subsystem_maps": [
    {
      "anagrp_id": "00103050h",
      "namespace": {
        "name": "/vol/vol1/namespacel",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "nsid": "00000001h"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target_nqn": "nqn.1992-01.example.com:string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
```

```
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#\_links]
[.api-collapsible-fifth-title]
_links
[#hosts]
[.api-collapsible-fifth-title]
hosts
[#default]
[.api-collapsible-fifth-title]
default

The default I/O queue parameters inherited by NVMe hosts in the NVMe
subsystem.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|count
|integer
a|The number of host I/O queue pairs.

|depth
|integer
a|The host I/O queue depth.
```

|===

```
[#io_queue]
[.api-collapsible-fifth-title]
io_queue
```

The properties of the submission queue used to submit I/O commands for execution by the NVMe controller.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|default
|link:#default[default]
a|The default I/O queue parameters inherited by NVMe hosts in the NVMe subsystem.
```

|===

```
[#namespace]
[.api-collapsible-fifth-title]
namespace
```

An NVMe namespace mapped to the NVMe subsystem.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the NVMe namespace.
```

```
|uuid
```

```
|string
a|The unique identifier of the NVMe namespace.
```

```
|===
```

```
[#subsystem_maps]
[.api-collapsible-fifth-title]
subsystem_maps
```

An NVMe namespace mapped to the NVMe subsystem.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|anagrpid
```

```
|string
```

```
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe
namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|namespace
```

```
|link:#namespace[namespace]
```

```
a|An NVMe namespace mapped to the NVMe subsystem.
```

```
|nsid
```

```
|string
```

```
a|The NVMe namespace identifier. This is an identifier used by an NVMe
controller to provide access to the NVMe namespace.
```

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_subsystem]
```

```
[.api-collapsible-fifth-title]
```

```
nvme_subsystem
```

An NVMe subsystem maintains configuration state and namespace access control for a set of NVMe-connected hosts.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|comment
```

```
|string
```

```
a|A configurable comment for the NVMe subsystem. Optional in POST and PATCH.
```

```
|delete_on_unmap
```

```
|boolean
```

```
a|An option that causes the subsystem to be deleted when the last subsystem map associated with it is deleted. This property defaults to _false_ when the subsystem is created.
```

```
|io_queue
|link:#io_queue[io_queue]
a|The properties of the submission queue used to submit I/O commands for
execution by the NVMe controller.
```

```
|serial_number
|string
a|The serial number of the NVMe subsystem.
```

```
|subsystem_maps
|array[link:#subsystem_maps[subsystem_maps]]
a|The NVMe namespaces mapped to the NVMe subsystem.
```

There is an added cost to retrieving property values for `subsystem_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|svm
|link:#svm[svm]
a|
```

```
|target_nqn
|string
a|The NVMe qualified name (NQN) used to identify the NVMe storage target.
```

```
|uuid
|string
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
```



```
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage NVMe namespaces
```

```
:leveloffset: +1
```

```
[[ID33fac6558cf7255a721a070b08f7c958]]
```

```
= Storage namespaces endpoint overview
```

```
== Overview
```

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

The NVMe namespace REST API allows you to create, update, delete and discover NVMe namespaces.

In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

```
== Performance monitoring
```

Performance of an NVMe namespace can be monitored by observing the ``metric.++`` and ``statistics.++`` properties. These properties show the performance of an NVMe namespace in terms of IOPS, latency, and throughput. The ``metric.++`` properties denote an average, whereas ``statistics.++`` properties denote a real-time monotonically increasing value aggregated across all nodes.

== Examples

=== Creating an NVMe namespace

This example creates a 300 gigabyte NVMe namespace, with 4096-byte blocks, in SVM `_svml_`, volume `_voll_`, configured for use by `_linux_` hosts. The ``return_records`` query parameter is used to retrieve properties of the newly created NVMe namespace in the POST response.

The API:

POST /api/storage/namespaces

The call:

```
curl -X POST 'https://<mgmt-  
ip>/api/storage/namespaces?return_records=true' -H 'accept:  
application/hal+json' -d '{ "svm": { "name": "svml" }, "os_type": "linux",  
"space": { "block_size": "4096", "size": "300G" }, "name" :  
"/vol/voll/namespacel" }'
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "uuid": "dccc3e6-cf4e-498f-bec6-f7897f945669",  
      "svm": {  
        "uuid": "6bf967fd-2a1c-11e9-b682-005056bbc17d",  
        "name": "svml",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/6bf967fd-2a1c-11e9-b682-005056bbc17d"  
          }  
        }  
      },  
      "name": "/vol/voll/namespacel",  
      "location": {
```

```

    "namespace": "namespace1",
    "volume": {
        "uuid": "71cd0dba-2a1c-11e9-b682-005056bbc17d",
        "name": "vol1",
        "_links": {
            "self": {
                "href": "/api/storage/volumes/71cd0dba-2a1c-11e9-b682-
005056bbc17d"
            }
        }
    },
    "enabled": true,
    "os_type": "linux",
    "space": {
        "block_size": 4096,
        "size": 322122547200,
        "used": 0,
        "guarantee": {
            "requested": false,
            "reserved": false
        }
    },
    "status": {
        "container_state": "online",
        "read_only": false,
        "state": "online"
    },
    "_links": {
        "self": {
            "href": "/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-
f7897f945669"
        }
    }
}
]
}
-----
'''

```

=== Updating an NVMe namespace

This example sets the `comment` property of an NVMe namespace.

```

# The API:
PATCH /api/storage/namespaces/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-f7897f945669' -H 'accept: application/hal+json' -d '{ "comment": "Data for the research department." }'
-----

'''

=== Retrieving NVMe namespaces

This example retrieves summary information for all online NVMe namespaces in SVM_svm1_. The `svm.name` and `status.state` query parameters are to find the desired NVMe namespaces.

-----

# The API:
GET /api/storage/namespaces

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/namespaces?svm.name=svm1&status.state=online' -H 'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "uuid": "5c254d22-96a6-42ac-aad8-0cd9ebd126b6",
      "svm": {
        "name": "svm1"
      },
      "name": "/vol/vol1/namespace2",
      "status": {
        "state": "online"
      },
      "_links": {
        "self": {
          "href": "/api/storage/namespaces/5c254d22-96a6-42ac-aad8-0cd9ebd126b6"
        }
      }
    }
  ]
}

```

```

    },
    {
      "uuid": "dccdc3e6-cf4e-498f-bec6-f7897f945669",
      "svm": {
        "name": "svm1"
      },
      "name": "/vol/vol1/namespace1",
      "status": {
        "state": "online"
      },
      "_links": {
        "self": {
          "href": "/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-
f7897f945669"
        }
      }
    },
    {
      "uuid": "be732687-20cf-47d2-a0e2-2a989d15661d",
      "svm": {
        "name": "svm1"
      },
      "name": "/vol/vol2/namespace3",
      "status": {
        "state": "online"
      },
      "_links": {
        "self": {
          "href": "/api/storage/namespaces/be732687-20cf-47d2-a0e2-
2a989d15661d"
        }
      }
    }
  ],
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/storage/namespaces?svm.name=svm1&status.state=online"
    }
  }
}
----

'''

```

=== Retrieving details for a specific NVMe namespace

In this example, the ``fields`` query parameter is used to request all fields, including advanced fields, that would not otherwise be returned by default for the NVMe namespace.

The API:

GET /api/storage/namespaces/{uuid}

The call:

curl -X GET 'https://<mgmt-ip>/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-f7897f945669?fields=**' -H 'accept: application/hal+json'

The response:

```
{
  "uuid": "dccdc3e6-cf4e-498f-bec6-f7897f945669",
  "svm": {
    "uuid": "6bf967fd-2a1c-11e9-b682-005056bbc17d",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/6bf967fd-2a1c-11e9-b682-005056bbc17d"
      }
    }
  },
  "name": "/vol/vol1/namespace1",
  "location": {
    "namespace": "namespace1",
    "volume": {
      "uuid": "71cd0dba-2a1c-11e9-b682-005056bbc17d",
      "name": "vol1",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/71cd0dba-2a1c-11e9-b682-005056bbc17d"
        }
      }
    }
  },
  "auto_delete": false,
  "enabled": true,
  "comment": "Data for the research department.",
  "os_type": "linux",
  "space": {
    "block_size": 4096,
```

```
"size": 322122547200,
"used": 0,
"guarantee": {
  "requested": false,
  "reserved": false
},
},
"status": {
  "container_state": "online",
  "mapped": true,
  "read_only": false,
  "state": "online"
},
"subsystem_map": {
  "nsid": "00000001h",
  "anagrp_id": "00000001h",
  "subsystem": {
    "uuid": "01f17d05-2be9-11e9-bed2-005056bbc17d",
    "name": "subsystem1",
    "_links": {
      "self": {
        "href": "/api/protocols/nvme/subsystems/01f17d05-2be9-11e9-bed2-005056bbc17d"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/protocols/nvme/subsystem-maps/dccdc3e6-cf4e-498f-bec6-f7897f945669/01f17d05-2be9-11e9-bed2-005056bbc17d"
    }
  }
},
"metric": {
  "timestamp": "2019-04-09T05:50:15Z",
  "duration": "PT15S",
  "status": "ok",
  "latency": {
    "other": 0,
    "total": 0,
    "read": 0,
    "write": 0
  },
  "iops": {
    "read": 0,
    "write": 0,
```



```

    "other": 0,
    "total": 0
  },
  "throughput": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"_links": {
  "self": {
    "href": "/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-
f7897f945669?fields=*"
  }
}
}
}
-----
'''

```

== Cloning NVMe namespaces

A clone of an NVMe namespace is an independent "copy" of the namespace that shares unchanged data blocks with the original. As blocks of the source and clone are modified, unique blocks are written for each. NVMe

namespace clones can be created quickly and consume very little space initially. They can be created for the purpose of back-up, or to replicate data for multiple consumers.

An NVMe namespace clone can also be set to auto-delete by setting the ``auto_delete`` property. If the namespace's volume is configured for automatic deletion, NVMe namespaces that have auto-delete enabled are deleted when a volume is nearly full to reclaim a target amount of free space in the volume.

=== Creating a new NVMe namespace clone

You create an NVMe namespace clone as you create any NVMe namespace -- a POST to

`xref:{relative_path}getting_started_with_the_ontap_rest_api.html#/NVMe/nvme_namespace_create[/storage/namespaces]`. Set ``clone.source.uuid`` or ``clone.source.name`` to identify the source NVMe namespace from which the clone is created. The NVMe namespace clone and its source must reside in the same volume.

The source NVMe namespace can reside in a Snapshot copy, in which case, the ``clone.source.name`` field must be used to identify it. Add ``/.snapshot/<snapshot_name>`` to the path after the volume name to identify the Snapshot copy. For example ``/vol/vol1/.snapshot/snap1/namespacel``.

The API:

POST /api/storage/namespaces

The call:

```
curl -X POST 'https://<mgmt-ip>/api/storage/namespaces' -H 'accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name":
"/vol/vol1/namespace2clone1", "clone": { "source": { "name":
"/vol/vol1/namespace2" } } }'
```

'''

=== Over-writing an existing NVMe namespace's data as a clone of another

You can over-write an existing NVMe namespace as a clone of another. You do this as a PATCH on the NVMe namespace to overwrite -- a PATCH to `xref:{relative_path}getting_started_with_the_ontap_rest_api.html#/NVMe/nvme_namespace_modify[/storage/namespaces/{uuid}]`. Set the ``clone.source.uuid`` or ``clone.source.name`` property to identify the source NVMe namespace from which the clone data is taken. The NVMe namespace

clone and its source must reside in the same volume.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH:

`auto_delete`, `subsystem_map`, `status.state`, and `uuid`.

The API:

PATCH /api/storage/namespaces/{uuid}

The call:

```
curl -X PATCH 'https://<mgmt-ip>/api/storage/namespaces/dccdc3e6-cf4e-498f-bec6-f7897f945669' -H 'accept: application/hal+json' -d '{ "clone": { "source": { "name": "/vol/vol1/namespace2" } } }'
```

'''

== Deleting an NVMe namespace

The API:

DELETE /api/storage/namespaces/{uuid}

The call:

```
curl -X DELETE 'https://<mgmt-ip>/api/storage/namespaces/5c254d22-96a6-42ac-aad8-0cd9ebd126b6' -H 'accept: application/hal+json'
```

'''

[[IDc97f7e9444233c95ee7926e2d2d9383d]]

= Retrieve NVMe namespaces

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/storage/namespaces`#

Introduced In: 9.6

Retrieves NVMe namespaces.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

- * `auto_delete`
- * `subsystem_map.+++`
- * `status.mapped`
- * `statistics.+++`
- * `metric.+++`

== Related ONTAP commands

- * `vserver nvme namespace show`
- * `vserver nvme subsystem map show`

== Learn more

* [xref:{relative_path}storage_namespaces_endpoint_overview.html\[DOC /storage/namespaces\]](#) to learn more and examples.

== Parameters

[cols=5*,options=header]
|===

Name	Type	In	Required	Description
------	------	----	----------	-------------

svm.uuid	string	query	False	a Filter by svm.uuid
----------	--------	-------	-------	----------------------

svm.name	string	query		
----------	--------	-------	--	--

```
|False
a|Filter by svm.name

|auto_delete
|boolean
|query
|False
a|Filter by auto_delete

|statistics.iops_raw.total
|integer
|query
|False
a|Filter by statistics.iops_raw.total

* Introduced in: 9.8

|statistics.iops_raw.read
|integer
|query
|False
a|Filter by statistics.iops_raw.read

* Introduced in: 9.8

|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other

* Introduced in: 9.8

|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write

* Introduced in: 9.8
```

```
|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total
```

* Introduced in: 9.8

```
|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read
```

* Introduced in: 9.8

```
|statistics.latency_raw.other
|integer
|query
|False
a|Filter by statistics.latency_raw.other
```

* Introduced in: 9.8

```
|statistics.latency_raw.write
|integer
|query
|False
a|Filter by statistics.latency_raw.write
```

* Introduced in: 9.8

```
|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp
```

* Introduced in: 9.8

```
|statistics.status
|string
```

```

|query
|False
a|Filter by statistics.status

* Introduced in: 9.8

|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write

* Introduced in: 9.8

|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read

* Introduced in: 9.8

|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total

* Introduced in: 9.8

|subsystem_map.anagrp_id
|string
|query
|False
a|Filter by subsystem_map.anagrp_id

|subsystem_map.nsid
|string
|query
|False
a|Filter by subsystem_map.nsid

```

```
|subsystem_map.subsystem.name
|string
|query
|False
a|Filter by subsystem_map.subsystem.name
```

```
|subsystem_map.subsystem.uuid
|string
|query
|False
a|Filter by subsystem_map.subsystem.uuid
```

```
|create_time
|string
|query
|False
a|Filter by create_time
```

* Introduced in: 9.7

```
|location.qtree.name
|string
|query
|False
a|Filter by location.qtree.name
```

```
|location.qtree.id
|integer
|query
|False
a|Filter by location.qtree.id
```

```
|location.volume.uuid
|string
|query
|False
a|Filter by location.volume.uuid
```

```
|location.volume.name
|string
```



```
|query  
|False  
a|Filter by location.volume.name
```

```
|location.namespace  
|string  
|query  
|False  
a|Filter by location.namespace
```

```
|enabled  
|boolean  
|query  
|False  
a|Filter by enabled
```

```
|status.mapped  
|boolean  
|query  
|False  
a|Filter by status.mapped
```

```
|status.read_only  
|boolean  
|query  
|False  
a|Filter by status.read_only
```

```
|status.container_state  
|string  
|query  
|False  
a|Filter by status.container_state
```

```
|status.state  
|string  
|query  
|False  
a|Filter by status.state
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|metric.latency.total
|integer
|query
|False
a|Filter by metric.latency.total
```

* Introduced in: 9.8

```
|metric.latency.read
|integer
|query
|False
a|Filter by metric.latency.read
```

* Introduced in: 9.8

```
|metric.latency.other
|integer
|query
|False
a|Filter by metric.latency.other
```

* Introduced in: 9.8

```
|metric.latency.write
|integer
|query
|False
a|Filter by metric.latency.write
```

* Introduced in: 9.8

```
|metric.status  
|string  
|query  
|False  
a|Filter by metric.status
```

* Introduced in: 9.8

```
|metric.iops.total  
|integer  
|query  
|False  
a|Filter by metric.iops.total
```

* Introduced in: 9.8

```
|metric.iops.read  
|integer  
|query  
|False  
a|Filter by metric.iops.read
```

* Introduced in: 9.8

```
|metric.iops.other  
|integer  
|query  
|False  
a|Filter by metric.iops.other
```

* Introduced in: 9.8

```
|metric.iops.write  
|integer  
|query  
|False  
a|Filter by metric.iops.write
```

* Introduced in: 9.8

```
|metric.duration
|string
|query
|False
a|Filter by metric.duration
```

* Introduced in: 9.8

```
|metric.timestamp
|string
|query
|False
a|Filter by metric.timestamp
```

* Introduced in: 9.8

```
|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write
```

* Introduced in: 9.8

```
|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read
```

* Introduced in: 9.8

```
|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total
```

* Introduced in: 9.8

```
|os_type
```

```
|string
|query
|False
a|Filter by os_type
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|space.guarantee.requested
|boolean
|query
|False
a|Filter by space.guarantee.requested
```

```
|space.guarantee.reserved
|boolean
|query
|False
a|Filter by space.guarantee.reserved
```

```
|space.used
|integer
|query
|False
a|Filter by space.used
```

```
|space.block_size
|integer
|query
|False
a|Filter by space.block_size
```

```
|space.size
|integer
|query
|False
a|Filter by space.size
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1
* Max value: 120
* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_namespace[nvme_namespace]]
a|

|===
```

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  ]
}
```

```

},
"comment": "string",
"create_time": "2018-06-04T19:00:00Z",
"location": {
  "namespace": "namespace1",
  "qtree": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": "1",
    "name": "qt1"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
},
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  }
}

```



```

    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "name": "/vol/volume1/mtree1/namespace1",
  "os_type": "string",
  "space": {
    "block_size": "512",
    "size": "1073741824",
    "used": 0
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    }
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"status": {
  "container_state": "string",
  "state": "online"
},
"subsystem_map": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "anagrpid": "00103050h",
  "nsid": "00000001h",
  "subsystem": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}

```

```

    },
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {

```

```

        "code": "string",
        "message": "string"
    }
],
"code": "4",
"message": "entry doesn't exist",
"target": "uuid"
}
}
=====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next

```

```

|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#source]
[.api-collapsible-fifth-title]
source

```

The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

```

[#clone]
[.api-collapsible-fifth-title]
clone

```

This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a

namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH:

`auto_delete` (unless specified in the request), `subsystem_map`,
`status.state`, and `uuid`.

```
[#qtree]
[.api-collapsible-fifth-title]
qtree
```

The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties `name` and `location.qtree.name` and/or
`location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.
```

```
|name
|string
a|The name of the qtree.
```

```
|===
```

```
[#volume]
[.api-collapsible-fifth-title]
```

volume

The volume in which the NVMe namespace is located. Valid in POST.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#location]
```

```
[.api-collapsible-fifth-title]
```

```
location
```

The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|namespace
```

```
|string
```

```
a|The base name component of the NVMe namespace. Valid in POST.
```

If properties ``name`` and ``location.namespace`` are specified in the same request, they must refer to the base name.

NVMe namespaces do not support rename.

```
|qtree
```

```
|link:#qtree[qtree]
```

```
a|The qtree in which the NVMe namespace is optionally located. Valid in POST.
```

If properties ``name`` and ``location.qtree.name`` and/or ``location.qtree.uuid`` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

```
|volume
```

```
|link:#volume[volume]
```

```
a|The volume in which the NVMe namespace is located. Valid in POST.
```

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

```
|===
```

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```

|write
|integer
a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

Performance numbers, such as IOPS latency and throughput

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized

```

failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#guarantee]

[.api-collapsible-fifth-title]

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|requested

|boolean

a|The requested space reservation policy for the NVMe namespace. If `_true_`, a space reservation is requested for the namespace; if `_false_`, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.

The space reservation policy for an NVMe namespace is determined by ONTAP.

* readOnly: 1
* Introduced in: 9.6

|reserved
|boolean
a|Reports if the NVMe namespace is space guaranteed.

This property is `_true_` if a space guarantee is requested and the containing volume and aggregate support the request. This property is `_false_` if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request.

|===

[#space]
[.api-collapsible-fifth-title]
space

The storage space related properties of the NVMe namespace.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|block_size
|integer

a|The size of blocks in the namespace in bytes.

Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.

Valid in POST.

|guarantee
|link:#guarantee[guarantee]

a|Properties that request and report the space guarantee for the NVMe namespace.

```
|size
|integer
a|The total provisioned size of the NVMe namespace.
```

NVMe namespaces do not support resize.

The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

```
* example: 1073741824
* Max value: 140737488355328
* Min value: 4096
* Introduced in: 9.6
* readCreate: 1
```

```
|used
|integer
a|The amount of space consumed by the main data stream of the NVMe
namespace.
```

This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used and an indicator for an out-of-space condition.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

```
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
```

```
[.api-collapsible-fifth-title]
```

```
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
```

```
|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.
```

```
|status
```



```
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#status]
[.api-collapsible-fifth-title]
status
```

Status information about the NVMe namespace.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|container_state
|string
a|The state of the volume and aggregate that contain the NVMe namespace.
Namespaces are only available when their containers are available.
```

```
|mapped
|boolean
a|Reports if the NVMe namespace is mapped to an NVMe subsystem.
```

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|read_only
|boolean
a|Reports if the NVMe namespace allows only read access.
```

```
|state
|string
a|The state of the NVMe namespace. Normal states for a namespace are
_online_ and _offline_. Other states indicate errors.
```

```
|===
```

```
[#subsystem]
[.api-collapsible-fifth-title]
subsystem
```

The NVMe subsystem to which the NVMe namespace is mapped.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
 |_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the NVMe subsystem.
```

```
|uuid
|string
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#subsystem_map]
[.api-collapsible-fifth-title]
subsystem_map
```

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|anagrpid
|string
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|nsid
|string
a|The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.
```

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe namespace is mapped.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_namespace]
[.api-collapsible-fifth-title]
nvme_namespace
```

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|auto_delete
```

```
|boolean
```

```
a|This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.
```

When set to `_true_`, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new NVMe namespace is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|comment
|string
a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time
|string
a|The time the NVMe namespace was created.

|enabled
|boolean
a|The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, you can check the `state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created.

|location
|link:#location[location]
a|The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

* Introduced in: 9.6
* readCreate: 1

|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput

|name
|string
a|The fully qualified path name of the NVMe namespace composed of a "/vol" prefix, the volume name, the (optional) qtree name and base name of the namespace. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

|os_type
|string

a|The operating system type of the NVMe namespace.

Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.

|space

|link:#space[space]

a|The storage space related properties of the NVMe namespace.

|statistics

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|status

|link:#status[status]

a|Status information about the NVMe namespace.

|subsystem_map

|link:#subsystem_map[subsystem_map]

a|The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The unique identifier of the NVMe namespace.

|===

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

```



```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[IDc3d00bb002acfffae970a0f4c8b6cf3d]]
= Create an NVMe namespace
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/storage/namespaces`#
```

Introduced In: 9.6

Creates an NVMe namespace.

== Required properties

- * `svm.uuid` or `svm.name` - Existing SVM in which to create the NVMe namespace.
- * `name`, `location.volume.name` or `location.volume.uuid` - Existing volume in which to create the NVMe namespace.
- * `name` or `location.namespace` - Base name for the NVMe namespace.
- * `os_type` - Operating system from which the NVMe namespace will be accessed. (Not used for clones, which are created based on the `os_type` of the source NVMe namespace.)
- * `space.size` - Size for the NVMe namespace. (Not used for clones, which are created based on the size of the source NVMe namespace.)

== Default property values

If not specified in POST, the following default property values are assigned:

- * `auto_delete` - `__false__`
- * `space.block_size` - `__4096__`

== Related ONTAP commands

```
* `volume file clone autodelete`  
* `volume file clone create`  
* `vserver nvme namespace create`
```

== Learn more

```
* xref:{relative_path}storage_namespaces_endpoint_overview.html[DOC  
/storage/namespaces]
```

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|return_records  
|boolean  
|query  
|False
```

```
a|The default is false. If set to true, the records are returned.
```

```
* Default value:
```

```
|===
```

== Request Body

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|auto_delete  
|boolean
```

```
a|This property marks the NVMe namespace for auto deletion when the volume  
containing the namespace runs out of space. This is most commonly set on  
namespace clones.
```

When set to `_true_`, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new NVMe namespace is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|clone

|link:#clone[clone]

a|This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH:
``auto_delete`` (unless specified in the request), ``subsystem_map``, ``status.state``, and ``uuid``.

|comment

|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time

|string

a|The time the NVMe namespace was created.

|enabled

|boolean

a|The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, you can check the ``state`` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created.

|location
|link:#location[location]
a|The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

* Introduced in: 9.6
* readCreate: 1

|name
|string
a|The fully qualified path name of the NVMe namespace composed of a "/vol" prefix, the volume name, the (optional) qtree name and base name of the namespace. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

|os_type
|string
a|The operating system type of the NVMe namespace.

Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.

|space
|link:#space[space]
a|The storage space related properties of the NVMe namespace.

|status
|link:#status[status]
a|Status information about the NVMe namespace.

|subsystem_map
|link:#subsystem_map[subsystem_map]
a|The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See

xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields[Requesting specific fields] to learn more.

|svm
|link:#svm[svm]
a|

|uuid
|string
a|The unique identifier of the NVMe namespace.

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "clone": {
    "source": {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "create_time": "2018-06-04T19:00:00Z",
  "location": {
    "namespace": "namespace1",
    "qtree": {
      "id": "1",
      "name": "qt1"
    },
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "name": "/vol/volume1/qtree1/namespace1",
  "os_type": "string",
  "space": {
    "block_size": "512",
    "size": "1073741824",
    "used": 0
  }
}
```

```

},
"status": {
  "container_state": "string",
  "state": "online"
},
"subsystem_map": {
  "anagrpid": "00103050h",
  "nsid": "000000001h",
  "subsystem": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#nvme_namespace[nvme_namespace]]
a|

|===

.Example response
[%collapsible%closed]
====

```

```
[source,json,subs=+macros]
{
  "records": [
    {
      "clone": {
        "source": {
          "name": "/vol/volume1/namespace1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "comment": "string",
      "create_time": "2018-06-04T19:00:00Z",
      "location": {
        "namespace": "namespace1",
        "qtree": {
          "id": "1",
          "name": "qt1"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      },
      "name": "/vol/volume1/qtree1/namespace1",
      "os_type": "string",
      "space": {
        "block_size": "512",
        "size": "1073741824",
        "used": 0
      },
      "status": {
        "container_state": "string",
        "state": "online"
      },
      "subsystem_map": {
        "anagrp_id": "00103050h",
        "nsid": "00000001h",
        "subsystem": {
          "name": "string",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
    }
  ]
}
```

```
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
]
}
====

== Error
```

Status: Default

ONTAP Error Response Codes

```
|====
| Error Code | Description

| 917927
| The specified volume was not found.

| 918236
| The specified `location.volume.uuid` and `location.volume.name` do not
refer to the same volume.

| 2621462
| The supplied SVM does not exist.

| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5242927
| The specified qtree was not found.

| 5242950
| The specified `location.qtree.id` and `location.qtree.name` do not refer
to the same qtree.

| 5374352
| An invalid name was provided for the NVMe namespace.

| 5374858
| The volume specified by `name` is not the same as that specified by
`location.volume`.

| 5374860
```



```

| The qtree specified by `name` is not the same as that specified by
`location.qtree`.

| 5374861
| The NVMe namespace base name specified by `name` is not the same as that
specified by `location.name`.

| 5374862
| No NVMe namespace path base name was provided for the namespace.

| 13565952
| The NVMe namespace clone request failed.

| 72089720
| NVMe namespaces cannot be created in Snapshot copies.

| 72089721
| The volume specified is in a load sharing mirror relationship.
Namespaces are not supported in load sharing mirrors.

| 72089722
| A negative size was provided for the NVMe namespace.

| 72089723
| The specified size is too small for the NVMe namespace.

| 72089724
| The specified size is too large for the NVMe namespace.

| 72089725
| A LUN or NVMe namespace already exists at the specified path.

| 72089727
| NVMe namespaces cannot be created on an SVM root volume.

| 72089728
| NVMe namespaces cannot be created on a FlexGroup volume.

| 72089732
| An NVMe namespace name can only contain characters A-Z, a-z, 0-9, "-",
".", "_", "{" and "}".

| 72090005
| The specified `clone.source.uuid` and `clone.source.name` do not refer
to the same NVMe namespace.

| 72090006

```

```

| The specified `clone.source` was not found.

| 72090007
| The specified `clone.source` was not found.

| 72090009
| An error occurred after successfully creating the NVMe namespace. Some
properties were not set.

| 72090012
| The property cannot be specified when creating an NVMe namespace clone.
The `target` property of the error object identifies the property.

| 72090013
| The property is required except when creating an NVMe namespace clone.
The `target` property of the error object identifies the property.

| 72090014
| No volume was specified for the NVMe namespace.

| 72090015
| An error occurred after successfully creating the NVMe namespace
preventing the retrieval of its properties.

| 72090033
| The `clone.source.uuid` property is not supported when specifying a
source NVMe namespace from a Snapshot copy.
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#source]
[.api-collapsible-fifth-title]
source
```

The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
```

a|The fully qualified path name of the clone source NVMe namespace composed of a "/vol" prefix, the volume name, the (optional) qtree name and base name of the namespace. Valid in POST and PATCH.

```
|uuid
|string
```

a|The unique identifier of the clone source NVMe namespace. Valid in POST and PATCH.

```
|===
```

```
[#clone]
```

[.api-collapsible-fifth-title]

clone

This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH:

`auto_delete` (unless specified in the request), `subsystem_map`,
`status.state`, and `uuid`.

[cols=3*,options=header]

|==

|Name

|Type

|Description

|source

|link:#source[source]

a|The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

|==

[#qtree]

[.api-collapsible-fifth-title]

qtree

The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties `name` and `location.qtree.name` and/or
`location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|id
```

```
|integer
```

a|The identifier for the qtree, unique within the qtree's volume.

```
|name
```

```
|string
```

a|The name of the qtree.

```
|===
```

```
[#volume]
```

```
[.api-collapsible-fifth-title]
```

volume

The volume in which the NVMe namespace is located. Valid in POST.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|The name of the volume.

```
|uuid
```

```
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|namespace
|string
```

```
a|The base name component of the NVMe namespace. Valid in POST.
```

If properties ``name`` and ``location.namespace`` are specified in the same request, they must refer to the base name.

NVMe namespaces do not support rename.

```
|qtree
```

```
|link:#qtree[qtree]
```

```
a|The qtree in which the NVMe namespace is optionally located. Valid in
POST.
```

If properties ``name`` and ``location.qtree.name`` and/or ``location.qtree.uuid`` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

|volume

|link:#volume[volume]

a|The volume in which the NVMe namespace is located. Valid in POST.

If properties `name` and `location.volume.name` and/or
`location.volume.uuid` are specified in the same request, they must refer
to the same volume.

NVMe namespaces do not support movement between volumes.

|===

[#iops]

[.api-collapsible-fifth-title]

iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

```
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```


The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

```
metric
```

Performance numbers, such as IOPS latency and throughput

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|duration
```

```
|string
```

```
a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#guarantee]
[.api-collapsible-fifth-title]
guarantee
```

Properties that request and report the space guarantee for the NVMe namespace.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|requested
```

```
|boolean
```

a|The requested space reservation policy for the NVMe namespace. If `_true_`, a space reservation is requested for the namespace; if `_false_`, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.

The space reservation policy for an NVMe namespace is determined by ONTAP.

* readOnly: 1

* Introduced in: 9.6

```
|reserved
```

```
|boolean
```

a|Reports if the NVMe namespace is space guaranteed.

This property is `_true_` if a space guarantee is requested and the containing volume and aggregate support the request. This property is `_false_` if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request.

```
|===
```

```
[#space]
```

```
[.api-collapsible-fifth-title]
```

space

The storage space related properties of the NVMe namespace.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|block_size
|integer
a|The size of blocks in the namespace in bytes.
```

Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.

Valid in POST.

```
|guarantee
|link:#guarantee[guarantee]
a|Properties that request and report the space guarantee for the NVMe namespace.
```

```
|size
|integer
a|The total provisioned size of the NVMe namespace.
```

NVMe namespaces do not support resize.

The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

```
* example: 1073741824
* Max value: 140737488355328
* Min value: 4096
* Introduced in: 9.6
* readCreate: 1
```

```
|used
|integer
a|The amount of space consumed by the main data stream of the NVMe namespace.
```

This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe

filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used and an indicator for an out-of-space condition.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

- * readOnly: 1
- * Introduced in: 9.6

|===

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#status]
[.api-collapsible-fifth-title]

status

Status information about the NVMe namespace.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|container_state

|string

a|The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available.

|mapped

|boolean

a|Reports if the NVMe namespace is mapped to an NVMe subsystem.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|read_only

|boolean

a|Reports if the NVMe namespace allows only read access.

|state

|string

a|The state of the NVMe namespace. Normal states for a namespace are `_online_` and `_offline_`. Other states indicate errors.

|===

[#subsystem]

[.api-collapsible-fifth-title]

subsystem

The NVMe subsystem to which the NVMe namespace is mapped.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the NVMe subsystem.
```

```
|uuid
|string
a|The unique identifier of the NVMe subsystem.
```

```
|===
```

```
[#subsystem_map]
[.api-collapsible-fifth-title]
subsystem_map
```

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|anagrpid
|string
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.
```

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed

by a lower case "h".

|nsid

|string

a|The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

|subsystem

|link:#subsystem[subsystem]

a|The NVMe subsystem to which the NVMe namespace is mapped.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#nvme_namespace]

[.api-collapsible-fifth-title]

nvme_namespace

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed, resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|auto_delete
```

```
|boolean
```

```
a|This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.
```

When set to `_true_`, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new NVMe namespace is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|clone
|link:#clone[clone]
a|This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH:

`auto_delete` (unless specified in the request), `subsystem_map`,
`status.state`, and `uuid`.

|comment
|string
a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time
|string
a|The time the NVMe namespace was created.

|enabled
|boolean
a|The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, you can check the `state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created.

|location
|link:#location[location]
a|The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

* Introduced in: 9.6
* readCreate: 1

|name
|string

a|The fully qualified path name of the NVMe namespace composed of a "/vol" prefix, the volume name, the (optional) qtree name and base name of the namespace. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

|os_type

|string

a|The operating system type of the NVMe namespace.

Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.

|space

|link:#space[space]

a|The storage space related properties of the NVMe namespace.

|status

|link:#status[status]

a|Status information about the NVMe namespace.

|subsystem_map

|link:#subsystem_map[subsystem_map]

a|The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The unique identifier of the NVMe namespace.

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[ID33766687dd6386e30115ca1957b4400b]]

= Delete an NVMe namespace

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/storage/namespaces/{uuid}`#

Introduced In: 9.6

Deletes an NVMe namespace.

== Related ONTAP commands

* `vserver nvme namespace delete`

== Learn more

* xref:{relative_path}storage_namespaces_endpoint_overview.html[DOC /storage/namespaces]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description


```

|uuid
|string
|path
|True
a|The unique identifier of the NVMe namespace to delete.

|allow_delete_while_mapped
|boolean
|query
|False
a|Allows deletion of a mapped NVMe namespace.
A mapped NVMe namespace might be in use. Deleting a mapped namespace also
deletes the namespace map and makes the data no longer available, possibly
causing a disruption in the availability of data.
*This parameter should be used with caution.*

* Default value:

|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

```

ONTAP Error Response Codes

|===
| Error Code | Description

| 72090006
| The specified namespace was not found.

| 72090007
| The specified namespace was not found.

| 72090016
| The namespace's aggregate is offline. The aggregate must be online to
modify or remove the namespace.

| 72090017

```

```

| The namespace's volume is offline. The volume must be online to modify
or remove the namespace.
|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]

```

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDc966b2f68729fb1e35c6a409f7e90571]]
= Retrieve an NVMe namespace

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/namespaces/{uuid}`#

*Introduced In:* 9.6

Retrieves an NVMe namespace.

== Expensive properties

There is an added cost to retrieving values for these properties. They are
not included by default in GET results and must be explicitly requested
using the `fields` query parameter. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `auto_delete`
* `subsystem_map.+++`
* `status.mapped`
* `statistics.+++`
* `metric.+++`

== Related ONTAP commands

* `vserver nvme namespace show`
* `vserver nvme subsystem map show`

== Learn more

* xref:{relative_path}storage_namespaces_endpoint_overview.html[DOC
/storage/namespaces]
```

```

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|The unique identifier of the NVMe namespace to retrieve.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|auto_delete
|boolean
a|This property marks the NVMe namespace for auto deletion when the volume
containing the namespace runs out of space. This is most commonly set on
namespace clones.

```

When set to `_true_`, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new NVMe namespace is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|comment

|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time

|string

a|The time the NVMe namespace was created.

|enabled

|boolean

a|The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, you can check the ``state`` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created.

|location

|link:#location[location]

a|The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

* Introduced in: 9.6

* readCreate: 1

|metric

|link:#metric[metric]

a|Performance numbers, such as IOPS latency and throughput

|name

|string

a|The fully qualified path name of the NVMe namespace composed of a "/vol" prefix, the volume name, the (optional) qtree name and base name of the namespace. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

|os_type

|string

a|The operating system type of the NVMe namespace.

Required in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.

|space

|link:#space[space]

a|The storage space related properties of the NVMe namespace.

|statistics

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|status

|link:#status[status]

a|Status information about the NVMe namespace.

|subsystem_map

|link:#subsystem_map[subsystem_map]

a|The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|svm
|link:#svm[svm]
a|

|uuid
|string
a|The unique identifier of the NVMe namespace.
```

```
|===
```

```
.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "create_time": "2018-06-04T19:00:00Z",
  "location": {
    "namespace": "namespace1",
    "qtree": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": "1",
      "name": "qt1"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
}
```



```

},
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"name": "/vol/volume1/mtree1/namespace1",
"os_type": "string",
"space": {
  "block_size": "512",
  "size": "1073741824",
  "used": 0
},
"statistics": {
  "iops_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput_raw": {

```

```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"status": {
    "container_state": "string",
    "state": "online"
},
"subsystem_map": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "anagrpid": "00103050h",
    "nsid": "000000001h",
    "subsystem": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 72090006

| The specified namespace was not found.

| 72090007

| The specified namespace was not found.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#source]
[.api-collapsible-fifth-title]
source

```

The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

```
[#clone]
[.api-collapsible-fifth-title]
clone
```

This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH:

`auto_delete` (unless specified in the request), `subsystem_map`,
`status.state`, and `uuid`.

```
[#qtree]
[.api-collapsible-fifth-title]
qtree
```

The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties `name` and `location.qtree.name` and/or
`location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|id
|integer
```

a|The identifier for the qtree, unique within the qtree's volume.

|name

|string

a|The name of the qtree.

|===

[#volume]

[.api-collapsible-fifth-title]

volume

The volume in which the NVMe namespace is located. Valid in POST.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the volume.

|uuid

|string

a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7

* Introduced in: 9.6

|===

[#location]
[.api-collapsible-fifth-title]
location

The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|namespace
|string
a|The base name component of the NVMe namespace. Valid in POST.

If properties `name` and `location.namespace` are specified in the same request, they must refer to the base name.

NVMe namespaces do not support rename.

|qtree
|link:#qtree[qtree]
a|The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

|volume
|link:#volume[volume]
a|The volume in which the NVMe namespace is located. Valid in POST.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

|===

[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===


```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS latency and throughput

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#guarantee]

[.api-collapsible-fifth-title]

guarantee

Properties that request and report the space guarantee for the NVMe namespace.

[cols=3*,options=header]

|===

|Name

```

|Type
|Description

|requested
|boolean
a|The requested space reservation policy for the NVMe namespace. If
_true_, a space reservation is requested for the namespace; if _false_,
the namespace is thin provisioned. Guaranteeing a space reservation
request for a namespace requires that the volume in which the namespace
resides also be space reserved and that the fractional reserve for the
volume be 100%.

The space reservation policy for an NVMe namespace is determined by ONTAP.

* readOnly: 1
* Introduced in: 9.6

|reserved
|boolean
a|Reports if the NVMe namespace is space guaranteed.

This property is _true_ if a space guarantee is requested and the
containing volume and aggregate support the request. This property is
_false_ if a space guarantee is not requested or if a space guarantee is
requested and either the containing volume and aggregate do not support
the request.

|===

[#space]
[.api-collapsible-fifth-title]
space

The storage space related properties of the NVMe namespace.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|block_size
|integer

```

a|The size of blocks in the namespace in bytes.

Valid in POST when creating an NVMe namespace that is not a clone of another. Disallowed in POST when creating a namespace clone.

Valid in POST.

|guarantee

|link:#guarantee[guarantee]

a|Properties that request and report the space guarantee for the NVMe namespace.

|size

|integer

a|The total provisioned size of the NVMe namespace.

NVMe namespaces do not support resize.

The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The maximum size is variable with respect to large NVMe namespace support in ONTAP. If large namespaces are supported, the maximum size is 128 TB (140737488355328 bytes) and if not supported, the maximum size is just under 16 TB (17557557870592 bytes). The minimum size supported is always 4096 bytes.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* example: 1073741824

* Max value: 140737488355328

* Min value: 4096

* Introduced in: 9.6

* readCreate: 1

|used

|integer

a|The amount of space consumed by the main data stream of the NVMe namespace.

This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge

of how the namespace blocks are utilized outside of ONTAP, this property should not be used and an indicator for an out-of-space condition.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* readOnly: 1
* Introduced in: 9.6

|===

[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|other
|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read
|integer

a|Performance metric for read I/O operations.

|total
|integer

a|Performance metric aggregated over all types of I/O operations.

|write
|integer

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]  
[.api-collapsible-fifth-title]  
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]  
[.api-collapsible-fifth-title]  
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```


a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw

|link:#latency_raw[latency_raw]

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#status]

[.api-collapsible-fifth-title]

status

Status information about the NVMe namespace.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|container_state
```

```
|string
```

a|The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available.

```
|mapped
```

```
|boolean
```

a|Reports if the NVMe namespace is mapped to an NVMe subsystem.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|read_only
```

```
|boolean
```

a|Reports if the NVMe namespace allows only read access.

```
|state
```

```
|string
```

a|The state of the NVMe namespace. Normal states for a namespace are `_online_` and `_offline_`. Other states indicate errors.

```
|===
```

```
[#subsystem]
```

```
[.api-collapsible-fifth-title]
```

subsystem

The NVMe subsystem to which the NVMe namespace is mapped.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the NVMe subsystem.

|uuid
|string
a|The unique identifier of the NVMe subsystem.

|===

```

```

[#subsystem_map]
[.api-collapsible-fifth-title]
subsystem_map

```

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|anagrpid

```

```

|string
a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe
namespace.

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed
by a lower case "h".

|nsid
|string
a|The NVMe namespace identifier. This is an identifier used by an NVMe
controller to provide access to the NVMe namespace.

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-
filled) followed by a lower case "h".

|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe namespace is mapped.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID32a671f57454887f0695e4f5a6f76a71]]
= Update an NVMe namespace

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/storage/namespaces/{uuid}`#

*Introduced In:* 9.6

Updates an NVMe namespace.

== Related ONTAP commands

* `volume file clone autodelete`
* `vserver nvme namespace modify`

== Learn more

* xref:{relative_path}storage_namespaces_endpoint_overview.html[DOC
/storage/namespaces]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type

```

```
| In
| Required
| Description
```

```
| uuid
| string
| path
| True
```

a|The unique identifier of the NVMe namespace to update.

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
| Name
| Type
| Description
```

```
| auto_delete
| boolean
```

a|This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.

When set to `_true_`, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new NVMe namespace is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
| clone
| link:#clone[clone]
```

a|This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing

namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH:

`auto_delete` (unless specified in the request), `subsystem_map`, `status.state`, and `uuid`.

|comment

|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time

|string

a|The time the NVMe namespace was created.

|enabled

|boolean

a|The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, you can check the `state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created.

|status

|link:#status[status]

a|Status information about the NVMe namespace.

|subsystem_map

|link:#subsystem_map[subsystem_map]

a|The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|svm


```

|link:#svm[svm]
a|

|uuid
|string
a|The unique identifier of the NVMe namespace.

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "clone": {
    "source": {
      "name": "/vol/volume1/namespace1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "create_time": "2018-06-04T19:00:00Z",
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "subsystem_map": {
    "anagrpid": "00103050h",
    "nsid": "00000001h",
    "subsystem": {
      "name": "string",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
=====

== Response

```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 13565952
```

```
| The namespace clone request failed.
```

```
| 72090005
```

```
| The specified `clone.source.uuid` and `clone.source.name` do not refer  
to the same LUN.
```

```
| 72090006
```

```
| The specified namespace was not found. This can apply to `clone.source`  
or the target namespace. The `target` property of the error object  
identifies the property.
```

```
| 72090007
```

```
| The specified namespace was not found. This can apply to `clone.source`  
or the target namespace. The `target` property of the error object  
identifies the property.
```

```
| 72090010
```

```
| An error occurred after successfully overwriting data for the namespace  
as a clone. Some properties were not modified.
```

```
| 72090011
```

```
| An error occurred after successfully modifying some of the properties of  
the namespace. Some properties were not modified.
```

```
| 72090016
```

```
| The namespace's aggregate is offline. The aggregate must be online to  
modify or remove the namespace.
```

```
| 72090017
```

```
| The namespace's volume is offline. The volume must be online to modify  
or remove the namespace.
```

```
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

====

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#source]
```

```
[.api-collapsible-fifth-title]
```

```
source
```

The source NVMe namespace for a namespace clone operation. This can be specified using property ``clone.source.uuid`` or ``clone.source.name``. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

|name
|string
a|The fully qualified path name of the clone source NVMe namespace composed of a "/vol" prefix, the volume name, the (optional) qtree name and base name of the namespace. Valid in POST and PATCH.

|uuid
|string
a|The unique identifier of the clone source NVMe namespace. Valid in POST and PATCH.

|===

[#clone]
[.api-collapsible-fifth-title]
clone

This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the patched namespace unless otherwise specified as part of the PATCH:

`auto_delete` (unless specified in the request), `subsystem_map`,
`status.state`, and `uuid`.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|source

|link:#source[source]

a|The source NVMe namespace for a namespace clone operation. This can be specified using property `clone.source.uuid` or `clone.source.name`. If both properties are supplied, they must refer to the same namespace.

Valid in POST to create a new NVMe namespace as a clone of the source.

Valid in PATCH to overwrite an existing NVMe namespace's data as a clone of another.

|===

```
[#qtree]
[.api-collapsible-fifth-title]
qtree
```

The qtree in which the NVMe namespace is optionally located. Valid in POST.

If properties ``name`` and ``location.qtree.name`` and/or ``location.qtree.uuid`` are specified in the same request, they must refer to the same qtree.

NVMe namespaces do not support rename.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.
```

```
|name
|string
a|The name of the qtree.
```

|===

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

The volume in which the NVMe namespace is located. Valid in POST.

If properties ``name`` and ``location.volume.name`` and/or

`location.volume.uuid` are specified in the same request, they must refer to the same volume.

NVMe namespaces do not support movement between volumes.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#location]
```

```
[.api-collapsible-fifth-title]
```

```
location
```

The location of the NVMe namespace within the ONTAP cluster. Valid in POST.

NVMe namespaces do not support rename, or movement between volumes.

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```



```
|===
```

```
[#metric]  
[.api-collapsible-fifth-title]  
metric
```

Performance numbers, such as IOPS latency and throughput

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|duration
```

```
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
```

```
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
```

```
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#guarantee]
[.api-collapsible-fifth-title]
guarantee
```

Properties that request and report the space guarantee for the NVMe namespace.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|requested
|boolean
```

a|The requested space reservation policy for the NVMe namespace. If `_true_`, a space reservation is requested for the namespace; if `_false_`, the namespace is thin provisioned. Guaranteeing a space reservation request for a namespace requires that the volume in which the namespace resides also be space reserved and that the fractional reserve for the volume be 100%.

The space reservation policy for an NVMe namespace is determined by ONTAP.

```
* readOnly: 1
* Introduced in: 9.6
```

```
|reserved
|boolean
a|Reports if the NVMe namespace is space guaranteed.
```

This property is `_true_` if a space guarantee is requested and the containing volume and aggregate support the request. This property is `_false_` if a space guarantee is not requested or if a space guarantee is requested and either the containing volume and aggregate do not support the request.

|===

[#space]

[.api-collapsible-fifth-title]

space

The storage space related properties of the NVMe namespace.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|guarantee

|link:#guarantee[guarantee]

a|Properties that request and report the space guarantee for the NVMe namespace.

|used

|integer

a|The amount of space consumed by the main data stream of the NVMe namespace.

This value is the total space consumed in the volume by the NVMe namespace, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways NVMe filesystems and applications utilize blocks within a namespace, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the namespace blocks are utilized outside of ONTAP, this property should not be used and an indicator for an out-of-space condition.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* readOnly: 1

* Introduced in: 9.6

|===

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

|===

```
[#latency_raw]
[.api-collapsible-fifth-title]
```

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

```

```

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

```

```

|write
|integer
a|Performance metric for write I/O operations.

```

```

|===

```

```

[#statistics]
[.api-collapsible-fifth-title]
statistics

```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

```

```

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

```

```

|latency_raw
|link:#latency_raw[latency_raw]

```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#status]

[.api-collapsible-fifth-title]

status

Status information about the NVMe namespace.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|container_state
|string
a|The state of the volume and aggregate that contain the NVMe namespace. Namespaces are only available when their containers are available.

|mapped
|boolean
a|Reports if the NVMe namespace is mapped to an NVMe subsystem.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|read_only
|boolean
a|Reports if the NVMe namespace allows only read access.

|state
|string
a|The state of the NVMe namespace. Normal states for a namespace are `_online_` and `_offline_`. Other states indicate errors.

|===

[#subsystem]
[.api-collapsible-fifth-title]
subsystem

The NVMe subsystem to which the NVMe namespace is mapped.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string

a|The name of the NVMe subsystem.

|uuid

|string

a|The unique identifier of the NVMe subsystem.

|===

[#subsystem_map]

[.api-collapsible-fifth-title]

subsystem_map

The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|anagrpid

|string

a|The Asymmetric Namespace Access Group ID (ANAGRPID) of the NVMe namespace.

The format for an ANAGRPID is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

|nsid

|string

a|The NVMe namespace identifier. This is an identifier used by an NVMe controller to provide access to the NVMe namespace.

The format for an NVMe namespace identifier is 8 hexadecimal digits (zero-filled) followed by a lower case "h".

```
|subsystem
|link:#subsystem[subsystem]
a|The NVMe subsystem to which the NVMe namespace is mapped.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#nvme_namespace]
[.api-collapsible-fifth-title]
nvme_namespace
```

An NVMe namespace is a collection of addressable logical blocks presented to hosts connected to the storage virtual machine using the NVMe over Fabrics protocol.

In ONTAP, an NVMe namespace is located within a volume. Optionally, it can be located within a qtree in a volume.

An NVMe namespace is created to a specified size using thin or thick provisioning as determined by the volume on which it is created. NVMe namespaces support being cloned. An NVMe namespace cannot be renamed,

resized, or moved to a different volume. NVMe namespaces do not support the assignment of a QoS policy for performance management, but a QoS policy can be assigned to the volume containing the namespace. See the NVMe namespace object model to learn more about each of the properties supported by the NVMe namespace REST API.

An NVMe namespace must be mapped to an NVMe subsystem to grant access to the subsystem's hosts. Hosts can then access the NVMe namespace and perform I/O using the NVMe over Fabrics protocol.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|auto_delete
```

```
|boolean
```

a|This property marks the NVMe namespace for auto deletion when the volume containing the namespace runs out of space. This is most commonly set on namespace clones.

When set to `_true_`, the NVMe namespace becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the namespace is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new NVMe namespace is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|clone
```

```
|link:#clone[clone]
```

a|This sub-object is used in POST to create a new NVMe namespace as a clone of an existing namespace, or PATCH to overwrite an existing namespace as a clone of another. Setting a property in this sub-object indicates that a namespace clone is desired.

When used in a PATCH, the patched NVMe namespace's data is over-written as a clone of the source and the following properties are preserved from the

patched namespace unless otherwise specified as part of the PATCH:

`auto_delete` (unless specified in the request), `subsystem_map`,
`status.state`, and `uuid`.

|comment

|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time

|string

a|The time the NVMe namespace was created.

|enabled

|boolean

a|The enabled state of the NVMe namespace. Certain error conditions cause the namespace to become disabled. If the namespace is disabled, you can check the `state` property to determine what error disabled the namespace. An NVMe namespace is enabled automatically when it is created.

|status

|link:#status[status]

a|Status information about the NVMe namespace.

|subsystem_map

|link:#subsystem_map[subsystem_map]

a|The NVMe subsystem with which the NVMe namespace is associated. A namespace can be mapped to zero (0) or one (1) subsystems.

There is an added cost to retrieving property values for `subsystem_map`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The unique identifier of the NVMe namespace.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[ID77c87afd5c83e8b0ef44bfb5f6eb072e]]
= Retrieve historical performance metrics for an NVMe namespace

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/namespaces/{uuid}/metrics`#
```

Introduced In: 9.8

Retrieves historical performance metrics for an NVMe namespace.

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|duration
|string
|query
|False
a|Filter by duration
```

```
|timestamp
|string
|query
|False
a|Filter by timestamp
```

```
|throughput.total
|integer
|query
|False
a|Filter by throughput.total
```

```
|throughput.read
|integer
|query
|False
a|Filter by throughput.read
```

```
|throughput.other
|integer
|query
|False
a|Filter by throughput.other
```

```
|throughput.write
|integer
|query
|False
a|Filter by throughput.write
```

```
|latency.total
|integer
|query
|False
a|Filter by latency.total
```

```
|latency.read
|integer
|query
```

```
|False  
a|Filter by latency.read
```

```
|latency.other  
|integer  
|query  
|False  
a|Filter by latency.other
```

```
|latency.write  
|integer  
|query  
|False  
a|Filter by latency.write
```

```
|status  
|string  
|query  
|False  
a|Filter by status
```

```
|iops.total  
|integer  
|query  
|False  
a|Filter by iops.total
```

```
|iops.read  
|integer  
|query  
|False  
a|Filter by iops.read
```

```
|iops.other  
|integer  
|query  
|False  
a|Filter by iops.other
```

```
|iops.write
```



```
|integer
|query
|False
a|Filter by iops.write
```

```
|uuid
|string
|path
|True
a|Unique identifier of the NVMe namespace.
```

```
|interval
|string
|query
|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:
```

```
* 1h: Metrics over the most recent hour sampled over 15 seconds.
* 1d: Metrics over the most recent day sampled over 5 minutes.
* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```
* Default value: 1
* Max value: 120
* Min value: 0
```

```
|fields
|array[string]
|query
|False
```

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|return_records

|boolean

|query

|False

a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of records

```

|records
|array[link:#records[records]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
    },
  ],
}

```

```

        "timestamp": "2017-01-25T11:20:13Z"
    }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions

```

```

[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
|===
```

```

|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#records]
[.api-collapsible-fifth-title]
records

Performance numbers, such as IOPS latency and throughput.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string

```


a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```
|===
```

```
//end collapsible .Definitions block  
====
```

```
:leveloffset: -1
```

```
:leveloffset: -1
```

```
= Name-services
```

```
:leveloffset: +1
```

```
[[ID1c28c3ebd6055529a3743da4b54df58c]]  
= Name-services overview
```

```
== Overview
```

ONTAP uses name-services to obtain information about users and clients. This information is used to authenticate users accessing data on or administering the storage system, and to map user credentials in a mixed environment. If the user database is stored in NIS or LDAP servers, NIS and LDAP name services need to be configured in ONTAP. DNS is used for resolving the hostnames. ns-switch is used to configure the SVMs with sources to search for network information and the order in which to search them.

```
= Manage DNS configurations
```

```
:leveloffset: +1
```

```
[[IDac5dcd19cb82a57fa0d0da707d4c431a]]
```

= Name-services dns endpoint overview

== Overview

Displays DNS information and controls the DNS subsystem. DNS domain name and DNS servers are required parameters.

== Retrieving DNS information

The DNS GET endpoint retrieves all of the DNS configurations for data SVMs.

DNS configuration for the cluster is retrieved via
xref:{relative_path}cluster-endpoint-overview.html[/api/cluster] .

== Examples

=== Retrieving all of the fields for all of the DNS configurations

The API:

/api/name-services/dns

The call:

```
curl -X GET "https://<mgmt-ip>/api/name-services/dns?fields=*" -H "accept: application/hal+json"
```

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
          }
        }
      },
      "domains": [
        "domainA.example.com"
      ],
      "servers": [
        "10.10.10.10"
      ]
    }
  ]
}
```

```

    ],
    "_links": {
      "self": {
        "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-005056b41bd1"
      }
    }
  },
  {
    "svm": {
      "uuid": "19076d35-6e27-11e8-b9b8-005056b41bd1",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/19076d35-6e27-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "domains": [
      "sample.example.com"
    ],
    "servers": [
      "11.11.11.11",
      "22.22.22.22",
      "33.33.33.33"
    ],
    "_links": {
      "self": {
        "href": "/api/name-services/dns/19076d35-6e27-11e8-b9b8-005056b41bd1"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/name-services/dns?fields=*"
  }
}
}
}
-----

```

=== Retrieving all DNS configurations whose domain name starts with _dom*_.

```

-----

# The API:
/api/name-services/dns

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/dns?domains=dom*" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
          }
        }
      },
      "domains": [
        "domainA.example.com"
      ],
      "_links": {
        "self": {
          "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    {
      "num_records": 1,
      "_links": {
        "self": {
          "href": "/api/name-services/dns?domains=dom*"
        }
      }
    }
  ]
}
-----

=== Retrieving the DNS configuration for a specific SVM

-----

```

```

# The API:
/api/name-services/dns/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/hal+json"

# The response:
{
  "svm": {
    "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
      }
    }
  },
  "domains": [
    "domainA.example.com"
  ],
  "servers": [
    "10.10.10.10"
  ],
  "_links": {
    "self": {
      "href": "/api/name-services/dns/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
}
-----

== Creating a DNS configuration

The DNS POST endpoint creates a DNS configuration for the specified SVM.

== Example

The following example shows a POST operation:

-----

# The API:
/api/name-services/dns

# The call:

```

```
curl -X POST "https://<mgmt-ip>/api/name-services/dns" -H "accept: application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": { \"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" }, \"domains\": [ \"domainA.example.com\" ], \"servers\": [ \"10.10.10.10\" ] }"
```

== Updating a DNS configuration

The DNS PATCH endpoint updates the DNS configuration for the specified SVM.

== Examples

=== Updating both the DNS domains and servers

The API:

/api/name-services/dns/{svm.uuid}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-b9b8-005056b41bd1" -H "accept: application/hal+json" -H "Content-Type: application/json" -d "{ \"domains\": [ \"domainA.example.com\", \"domainB.example.com\" ], \"servers\": [ \"10.10.10.10\", \"10.10.10.11\" ] }"
```

=== Updating the DNS servers only

The API:

/api/name-services/dns/{svm.uuid}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-b9b8-005056b41bd1" -H "accept: application/hal+json" -H "Content-Type: application/json" -d "{ \"servers\": [ \"10.10.10.10\" ] }"
```

== Deleting a DNS configuration

The DNS DELETE endpoint deletes the DNS configuration for the specified SVM.

== Example

The following example shows a DELETE operation.

The API:

/api/name-services/dns/{svm.uuid}

The call:

curl -X DELETE "https://<mgmt-ip>/api/name-services/dns/179d3c85-7053-11e8-b9b8-005056b41bd1" -H "accept: application/hal+json"

[[IDb47036b61ff81609c645df8fb2fab231]]

= Retrieve DNS configurations for all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/name-services/dns`#

Introduced In: 9.6

Retrieves the DNS configurations of all data SVMs.

DNS configuration for the cluster is retrieved and managed via
xref:{relative_path}cluster-endpoint-overview.html[/api/cluster] .

== Related ONTAP commands

* `vserver services name-service dns show`
* `vserver services name-service dns check`

== Learn more

* xref:{relative_path}name-services_dns_endpoint_overview.html[DOC /name-services/dns]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

```
|In
|Required
|Description

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|servers
|string
|query
|False
a|Filter by servers

|domains
|string
|query
|False
a|Filter by domains

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
```

```

|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.


|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|num_records
|integer
a|Number of DNS domain records
```

```
|records
|array[link:#dns[dns]]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "domains": [
        "example.com",
        "example2.example3.com"
      ],
      "servers": [
        "10.224.65.20",
        "2001:db08:a0b:12f0::1"
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    }
  ],
}
```

```

      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#dns]
[.api-collapsible-fifth-title]
dns

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|domains
|array[string]
a|A list of DNS domains.
Domain names have the following requirements:

* The name must contain only the following characters: A through Z,
a through z, 0 through 9, ".", "-" or "_".
* The first character of each label, delimited by ".", must be one
of the following characters: A through Z or a through z or 0
through 9.
* The last character of each label, delimited by ".", must be one of
the following characters: A through Z, a through z, or 0 through 9.
* The top level domain must contain only the following characters: A
through Z, a through z.
* The system reserves the following names:"all", "local", and "localhost".

|servers
|array[string]
a|The list of IP addresses of the DNS servers. Addresses can be either
IPv4 or IPv6 addresses.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

```

====

[[ID138bbc77add116e487577ce62af8f85f]]

= Create DNS domain and server configurations

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/name-services/dns`#

Introduced In: 9.6

Creates DNS domain and server configurations for an SVM.

== Important notes

- * Each SVM can have only one DNS configuration.
- * The domain name and the servers fields cannot be empty.
- * IPv6 must be enabled if IPv6 family addresses are specified in the `servers` field.
- * Configuring more than one DNS server is recommended to avoid a single point of failure.
- * The DNS server specified using the `servers` field is validated during this operation.

The validation fails in the following scenarios:

- . The server is not a DNS server.
- . The server does not exist.
- . The server is unreachable.

== Parameters

[cols=5*,options=header]

|====

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|domains

|array[string]

a|A list of DNS domains.

Domain names have the following requirements:

- * The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-" or "_".

- * The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.

- * The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.

- * The top level domain must contain only the following characters: A through Z, a through z.

- * The system reserves the following names:"all", "local", and "localhost".

|servers

|array[string]

a|The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.

|svm

|link:#svm[svm]

a|

|===

.Example request

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "domains": [
    "example.com",
    "example2.example3.com"
  ],
  "servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response
```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of DNS domain records


|records
|array[link:#dns[dns]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "domains": [
        "example.com",
        "example2.example3.com"
      ],
      "servers": [
        "10.224.65.20",
        "2001:db08:a0b:12f0::1"
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 2621706

| The specified SVM UUID is incorrect for the specified SVM name

| 8847360

| Only admin or data SVMs allowed

| 8847361

| Exceeded the maximum number of domains allowed. Maximum of six domains only

| 8847362

| Exceeded the maximum number of name servers allowed. Maximum of three name servers only

| 8847392

| Domain name cannot be an IP address

| 8847393

| Top level domain name is invalid

| 8847399

| One or more of the specified DNS servers do not exist or cannot be reached

| 8847394

| FQDN name violated the limitations

| 9240587

| FQDN name cannot be empty

| 9240588

| FQDN name is too long. Maximum supported length: 255 characters

| 9240590

| FQDN name is reserved. Following names are reserved: "all", "local" and "localhost"

| 9240607

| One of the FQDN labels is too long. Maximum supported length: 63 characters

```

| 13434916
| The SVM is in the process of being created. Wait a few minutes, and then
try the command again.

| 23724130
| Cannot use an IPv6 name server address because there are no IPv6 LIFs
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name

```

```
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#dns]
[.api-collapsible-fifth-title]
dns
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|domains
|array[string]
a|A list of DNS domains.
Domain names have the following requirements:
```

- * The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".
- * The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.
- * The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.
- * The top level domain must contain only the following characters: A through Z, a through z.
- * The system reserves the following names:"all", "local", and "localhost".

```
|servers
|array[string]
a|The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
```

```
|svm
|link:#svm[svm]
```



```

a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[IDa4ffb8a5ddbe1d5f3e01b11f79ff051b]]
= Delete a DNS domain configuration
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/name-services/dns/{svm.uuid}`#
```

Introduced In: 9.6

Deletes DNS domain configuration of the specified SVM.

== Related ONTAP commands

* `vserver services name-service dns delete`

== Learn more

* xref:{relative_path}name-services_dns_endpoint_overview.html[DOC /name-
services/dns]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
```

```
|Required
|Description

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
```

```

    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

```

```

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|arguments

```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[IDe3ed499c07d5255d8f332db48676c2f7]]
```

```
= Retrieve DNS domain and server configurations
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/name-services/dns/{svm.uuid}`#
```

```
*Introduced In:* 9.6
```

Retrieves DNS domain and server configuration of an SVM. By default, both DNS domains and servers are displayed.

DNS configuration for the cluster is retrieved and managed via `xref:{relative_path}cluster-endpoint-overview.html[/api/cluster]` .

```
== Related ONTAP commands
```

```
* `vserver services name-service dns show`
```

```
* `vserver services name-service dns check`
```

```
== Learn more
```

```
* xref:{relative_path}name-services_dns_endpoint_overview.html[DOC /name-services/dns]
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the SVM to which this object belongs.
```

```
|fields
```

```
|array[string]
```

```
|query
```

```
|False
```

```
a|Specify the fields to return.
```

```
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|domains
```

```
|array[string]
```

a|A list of DNS domains.

Domain names have the following requirements:

- * The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".
- * The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.
- * The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.
- * The top level domain must contain only the following characters: A through Z, a through z.
- * The system reserves the following names:"all", "local", and "localhost".

|servers

|array[string]

a|The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.

|svm

|link:#svm[svm]

a|

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "domains": [
    "example.com",
    "example2.example3.com"
  ],
  "servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ],
```

```

"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```



```

}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID8d167b72c3511bcf6d2c2eae403c02d0]]
= Update DNS domain and server configurations

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/name-services/dns/{svm.uuid}`#

*Introduced In:* 9.6

Updates DNS domain and server configurations of an SVM.

```

== Important notes

- * Both DNS domains and servers can be modified.
- * The domains and servers fields cannot be empty.
- * IPv6 must be enabled if IPv6 family addresses are specified for the `servers` field.
- * The DNS server specified using the `servers` field is validated during this operation.

The validation fails in the following scenarios:

- . The server is not a DNS server.
- . The server does not exist.
- . The server is unreachable.

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the SVM to which this object belongs.
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|domains
```

```
|array[string]
```

```
a|A list of DNS domains.
```

```
Domain names have the following requirements:
```

```
* The name must contain only the following characters: A through Z,  
a through z, 0 through 9, ".", "-" or "_".
```

```
* The first character of each label, delimited by ".", must be one  
of the following characters: A through Z or a through z or 0  
through 9.
```

```
* The last character of each label, delimited by ".", must be one of  
the following characters: A through Z, a through z, or 0 through 9.
```

```
* The top level domain must contain only the following characters: A  
through Z, a through z.
```

```
* The system reserves the following names:"all", "local", and "localhost".
```

```
|servers
```

```
|array[string]
```

```
a|The list of IP addresses of the DNS servers. Addresses can be either  
IPv4 or IPv6 addresses.
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

```
.Example request
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{  
  "domains": [  
    "example.com",  
    "example2.example3.com"  
  ],  
  "servers": [  
    "10.224.65.20",  
    "2001:db08:a0b:12f0::1"  
  ],  
  "svm": {  
    "name": "svm1",  
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"  
  }  
}
```

```
}  
====
```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

```
|===  
| Error Code | Description  
  
| 8847360  
| Only admin or data SVMs allowed  
  
| 8847361  
| Exceeded the maximum number of domains allowed. Maximum of six domains  
only  
  
| 8847362  
| Exceeded the maximum number of name servers allowed. Maximum of three  
name servers only  
  
| 8847392  
| Domain name cannot be an IP address  
  
| 8847393  
| Top level domain name is invalid  
  
| 8847394  
| FQDN name violated the limitations  
  
| 8847399  
| One or more of the specified DNS servers do not exist or cannot be  
reached  
  
| 9240587  
| FQDN name cannot be empty  
  
| 9240588  
| FQDN name is too long. Maximum supported length: 255 characters
```

```
| 9240590
| FQDN name is reserved. Following names are reserved: "all", "local" and
"localhost"
```

```
| 9240607
| One of the FQDN labels is too long. Maximum supported length: 63
characters
```

```
| 23724130
| Cannot use an IPv6 name server address because there are no IPv6 LIFs
|===
```

```
== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description

|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#dns]
[.api-collapsible-fifth-title]
dns
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|domains
|array[string]
a|A list of DNS domains.
```

Domain names have the following requirements:

- * The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-" or "_".
- * The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.
- * The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.
- * The top level domain must contain only the following characters: A through Z, a through z.
- * The system reserves the following names: "all", "local", and "localhost".

```
|servers
|array[string]
a|The list of IP addresses of the DNS servers. Addresses can be either
```


IPv4 or IPv6 addresses.

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```

|===

```

```

//end collapsible .Definitions block
====

```

```

:leveloffset: -1

```

```

= Manage LDAP server configurations

```

```

:leveloffset: +1

```

```

[[ID8fc141bfae68776b944b05401616f551]]
= Name-services LDAP endpoint overview

```

```

== Overview

```

LDAP servers are used to centrally maintain user information. LDAP configurations must be set up to lookup information stored in the LDAP directory on the external LDAP servers. This API is used to retrieve and manage LDAP server configurations.

```

== Retrieving LDAP information

```

The LDAP GET endpoint retrieves all of the LDAP configurations in the cluster.

== Examples

=== Retrieving all of the fields for all LDAP configurations

'''

The API:

/api/name-services/ldap

The call:

```
curl -X GET "https://<mgmt-ip>/api/name-services/ldap?fields=*" -H
"accept: application/hal+json"
```

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
          }
        }
      },
      "servers": [
        "10.10.10.10",
        "domainB.example.com"
      ],
      "schema": "ad_idmu",
      "port": 389,
      "min_bind_level": "anonymous",
      "bind_dn": "cn=Administrators,cn=users,dc=domainA,dc=example,dc=com",
      "base_dn": "dc=domainA,dc=example,dc=com",
      "base_scope": "subtree",
      "use_start_tls": true,
      "session_security": "none",
      "_links": {
        "self": {
          "href": "/api/name-services/ldap/179d3c85-7053-11e8-b9b8-
```

```

005056b41bd1"
    }
  },
  {
    "svm": {
      "uuid": "6a52023b-7066-11e8-b9b8-005056b41bd1",
      "name": "vs2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/6a52023b-7066-11e8-b9b8-005056b41bd1"
        }
      }
    },
    "servers": [
      "11.11.11.11"
    ],
    "schema": "rfc_2307",
    "port": 389,
    "min_bind_level": "simple",
    "bind_dn": "cn=Administrators,cn=users,dc=domainB,dc=example,dc=com",
    "base_dn": "dc=domainB,dc=example,dc=com",
    "base_scope": "subtree",
    "use_start_tls": true,
    "session_security": "sign",
    "_links": {
      "self": {
        "href": "/api/name-services/ldap/6a52023b-7066-11e8-b9b8-
005056b41bd1"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/name-services/ldap?fields=*"
  }
}
}
}
----

'''

```

=== Retrieving all of the LDAP configurations that have the
 _use_start_tls_ set to _true_

```

'''

----

# The API:
/api/name-services/ldap

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap?use_start_tls=true"
-H "accept: application/hal+json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "6a52023b-7066-11e8-b9b8-005056b41bd1",
        "name": "vs2",
        "_links": {
          "self": {
            "href": "/api/svm/svms/6a52023b-7066-11e8-b9b8-005056b41bd1"
          }
        }
      },
      "use_start_tls": true,
      "_links": {
        "self": {
          "href": "/api/name-services/ldap/6a52023b-7066-11e8-b9b8-005056b41bd1"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/name-services/ldap?use_start_tls=true"
    }
  }
}

----

'''

=== Retrieving the LDAP configuration of a specific SVM

```

```

'''

----

# The API:
/api/name-services/ldap/{svm.uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/name-services/ldap/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/json"

# The response:
{
  "svm": {
    "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
      }
    }
  },
  "servers": [
    "10.10.10.10",
    "domainB.example.com"
  ],
  "schema": "ad_idmu",
  "port": 389,
  "min_bind_level": "anonymous",
  "bind_dn": "cn=Administrators,cn=users,dc=domainA,dc=example,dc=com",
  "base_dn": "dc=domainA,dc=example,dc=com",
  "base_scope": "subtree",
  "use_start_tls": true,
  "session_security": "none",
  "_links": {
    "self": {
      "href": "/api/name-services/ldap/179d3c85-7053-11e8-b9b8-005056b41bd1"
    }
  }
}

----

'''

== Creating an LDAP configuration

```

The LDAP POST endpoint creates an LDAP configuration for the specified SVM.

== Examples

=== Creating an LDAP configuration with all the fields specified

'''

The API:

/api/name-services/ldap

The call:

```
curl -X POST "https://<mgmt-ip>/api/name-services/ldap" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" }, \"servers\": [
\"10.10.10.10\", \"domainB.example.com\" ], \"schema\": \"ad_idmu\",
\"port\": 389, \"min_bind_level\": \"anonymous\", \"bind_dn\":
\"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\",
\"base_scope\": \"subtree\", \"use_start_tls\": false,
\"session_security\": \"none\"}"
```

'''

=== Creating an LDAP configuration with Active Directory domain and preferred Active Directory servers specified

'''

The API:

/api/name-services/ldap

The call:

```
curl -X POST "https://<mgmt-ip>/api/name-services/ldap" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"name\": \"vs2\" }, \"ad_domain\": \"domainA.example.com\",
\"preferred_ad_servers\": [ \"11.11.11.11\" ], \"port\": 389, \"bind_dn\":
\"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\",
\"session_security\": \"none\"}"
```

'''

== Creating an LDAP configuration with a number of optional fields not specified

'''

The API:
/api/name-services/ldap

The call:
curl -X POST "https://<mgmt-ip>/api/name-services/ldap" -H "accept: application/hal+json" -H "Content-Type: application/json" -d "{ \"svm\": { \"name\": \"vs2\" }, \"servers\": [\"11.11.11.11\"], \"port\": 389, \"bind_dn\": \"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\", \"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\", \"session_security\": \"none\"}"

'''

== Updating an LDAP configuration

The LDAP PATCH endpoint updates the LDAP configuration for the specified SVM. The following example shows a PATCH operation:

The API:
/api/name-services/ldap/{svm.uuid}

The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/ldap/179d3c85-7053-11e8-b9b8-005056b41bd1" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"servers\": [\"55.55.55.55\"], \"schema\": \"ad_idmu\", \"port\": 636, \"use_start_tls\": false }"

'''

== Deleting an LDAP configuration

The LDAP DELETE endpoint deletes the LDAP configuration for the specified SVM. The following example shows a DELETE operation:

The API:

/api/name-services/ldap/{svm.uuid}

The call:

curl -X DELETE "https://<mgmt-ip>/api/name-services/ldap/179d3c85-7053-11e8-b9b8-005056b41bd1" -H "accept: application/hal+json"

'''

[[ID7fd1e84597bddd4e8349058adf2e48b3]]

= Retrieve an LDAP configuration for all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/name-services/ldap`#

Introduced In: 9.6

Retrieves the LDAP configurations for all SVMs.

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|schema

|string

|query

|False

a|Filter by schema

* Introduced in: 9.7

```
|session_security
|string
|query
|False
a|Filter by session_security
```

* Introduced in: 9.7

```
|base_dn
|string
|query
|False
a|Filter by base_dn
```

* Introduced in: 9.7

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

* Introduced in: 9.7

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

* Introduced in: 9.7

```
|port
|integer
|query
|False
a|Filter by port
```

* Introduced in: 9.7

```
|min_bind_level
```

```
|string
|query
|False
a|Filter by min_bind_level
```

* Introduced in: 9.7

```
|use_start_tls
|boolean
|query
|False
a|Filter by use_start_tls
```

* Introduced in: 9.7

```
|preferred_ad_servers
|string
|query
|False
a|Filter by preferred_ad_servers
```

* Introduced in: 9.7

```
|servers
|string
|query
|False
a|Filter by servers
```

* Introduced in: 9.7

```
|base_scope
|string
|query
|False
a|Filter by base_scope
```

* Introduced in: 9.7

```
|bind_dn
|string
|query
```

```

|False
a|Filter by bind_dn

* Introduced in: 9.7


|ad_domain
|string
|query
|False
a|Filter by ad_domain

* Introduced in: 9.7


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

```

```
* Default value: 1
* Max value: 120
* Min value: 0
```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of LDAP records.
```

```
|records
|array[link:#ldap_service[ldap_service]]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```

    "_links": {
      "next": {
        "href": "/api/resourcelink"
      },
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "records": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ad_domain": "string",
        "base_dn": "string",
        "base_scope": "string",
        "bind_dn": "string",
        "min_bind_level": "string",
        "port": "389",
        "preferred_ad_servers": [
          "string"
        ],
        "schema": "string",
        "servers": [
          "string"
        ],
        "session_security": "string",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
      }
    ]
  }
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

```



```
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#ldap_service]
[.api-collapsible-fifth-title]
ldap_service
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|ad_domain
|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST and PATCH.
```

```

|base_dn
|string
a|Specifies the default base DN for all searches.

|base_scope
|string
a|Specifies the default search scope for LDAP queries:

* base - search the named entry only
* onelevel - search all entries immediately below the DN
* subtree - search the named DN entry and the entire subtree below the DN

|bind_dn
|string
a|Specifies the user that binds to the LDAP servers.

|min_bind_level
|string
a|The minimum bind authentication level. Possible values are:

* anonymous - anonymous bind
* simple - simple bind
* sasl - Simple Authentication and Security Layer (SASL) bind

|port
|integer
a|The port used to connect to the LDAP Servers.

|preferred_ad_servers
|array[string]
a|

|schema
|string
a|The name of the schema template used by the SVM.

* AD-IDMU - Active Directory Identity Management for UNIX
* AD-SFU - Active Directory Services for UNIX
* MS-AD-BIS - Active Directory Identity Management for UNIX
* RFC-2307 - Schema based on RFC 2307

```

* Custom schema

|servers

|array[string]

a|

|session_security

|string

a|Specifies the level of security to be used for LDAP communications:

* none - no signing or sealing

* sign - sign LDAP traffic

* seal - seal and sign LDAP traffic

|svm

|link:#svm[svm]

a|

|use_start_tls

|boolean

a|Specifies whether or not to use Start TLS over LDAP connections.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID00e197cebc7cd006fc6ce75868291920]]
= Create an LDAP configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/name-services/ldap`#

```

Introduced In: 9.6

Creates an LDAP configuration for an SVM.

== Important notes

- * Each SVM can have one LDAP configuration.
- * The LDAP servers and Active Directory domain are mutually exclusive fields. These fields cannot be empty. At any point in time, either the LDAP servers or Active Directory domain must be populated.
- * LDAP configuration with Active Directory domain cannot be created on an admin SVM.
- * IPv6 must be enabled if IPv6 family addresses are specified.

=== The following parameters are optional:

- * preferred AD servers
- * schema
- * port
- * min_bind_level
- * bind_password
- * base_scope
- * use_start_tls
- * session_security

Configuring more than one LDAP server is recommended to avoid a single point of failure.

Both FQDNs and IP addresses are supported for the "servers" field.

The Active Directory domain or LDAP servers are validated as part of this operation.

LDAP validation fails in the following scenarios:

- . The server does not have LDAP installed.
- . The server or Active Directory domain is invalid.
- . The server or Active Directory domain is unreachable.

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

```

|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_domain
|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST and PATCH.

|base_dn
|string
a|Specifies the default base DN for all searches.

|base_scope
|string
a|Specifies the default search scope for LDAP queries:

* base - search the named entry only
* onelevel - search all entries immediately below the DN
* subtree - search the named DN entry and the entire subtree below the DN

|bind_dn
|string
a|Specifies the user that binds to the LDAP servers.

```

```

|bind_password
|string
a|Specifies the bind password for the LDAP servers.

|min_bind_level
|string
a|The minimum bind authentication level. Possible values are:

* anonymous - anonymous bind
* simple - simple bind
* sasl - Simple Authentication and Security Layer (SASL) bind

|port
|integer
a|The port used to connect to the LDAP Servers.

|preferred_ad_servers
|array[string]
a|

|schema
|string
a|The name of the schema template used by the SVM.

* AD-IDMU - Active Directory Identity Management for UNIX
* AD-SFU - Active Directory Services for UNIX
* MS-AD-BIS - Active Directory Identity Management for UNIX
* RFC-2307 - Schema based on RFC 2307
* Custom schema

|servers
|array[string]
a|

|session_security
|string
a|Specifies the level of security to be used for LDAP communications:

* none - no signing or sealing
* sign - sign LDAP traffic
* seal - seal and sign LDAP traffic

```

```
|svm
|link:#svm[svm]
a|

|use_start_tls
|boolean
a|Specifies whether or not to use Start TLS over LDAP connections.
```

```
|===
```

```
.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "ad_domain": "string",
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "bind_password": "string",
  "min_bind_level": "string",
  "port": "389",
  "preferred_ad_servers": [
    "string"
  ],
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

== Response
```

Status: 201, Created

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|num_records
|integer
a|Number of LDAP records.

|records
|array[link:#ldap_service[ldap_service]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "ad_domain": "string",
      "base_dn": "string",
      "base_scope": "string",
      "bind_dn": "string",
      "bind_password": "string",
      "min_bind_level": "string",
      "port": "389",
      "preferred_ad_servers": [
        "string"
      ],
      "schema": "string",
      "servers": [
        "string"
      ],
      "session_security": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}

```

====

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

262186	
--------	--

	LDAP Servers cannot be used with Active Directory domain and/or preferred Active Directory servers
--	----------------------------------------------------------------------------------------------------

2621488	
---------	--

	Invalid SVM context
--	---------------------

2621706	
---------	--

	The specified SVM UUID is incorrect for the specified SVM name
--	----------------------------------------------------------------

4915203	
---------	--

	The specified LDAP schema does not exist
--	------------------------------------------

4915207	
---------	--

	The specified LDAP servers or preferred Active Directory servers contain duplicate server entries
--	---------------------------------------------------------------------------------------------------

4915229	
---------	--

	DNS resolution failed due to an internal error. Contact technical support if this issue persists
--	--------------------------------------------------------------------------------------------------

4915231	
---------	--

	DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured
--	-------------------------------------------------------------------------------------------------------------------

23724132	
----------	--

	DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured
--	--------------------------------------------------------------------------------------------------------

4915234	
---------	--

	The specified LDAP server or preferred Active Directory server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast
--	----------------------------------------------------------------------------------------------------------------------------------------------------------------

4915248	
---------	--

| LDAP servers cannot be empty or "-". Specified Active Directory domain is invalid because it is empty or "-" or it contains either the special characters or "-" at the start or end of the domain)

| 4915251

| STARTTLS and LDAPS cannot be used together

| 4915257

| The LDAP configuration is invalid. Verify that bind-dn and bind password are correct

| 4915258

| The LDAP configuration is invalid. Verify that the Active Directory domain or servers are reachable and that the network configuration is correct

| 4915259

| LDAP configurations with Active Directory domains are not supported on admin SVM.

| 13434916

| The SVM is in the process of being created. Wait a few minutes, and then try the command again.

| 23724130

| Cannot use an IPv6 name server address because there are no IPv6 LIFs

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#href]

[.api-collapsible-fifth-title]

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

```

|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#ldap_service]
[.api-collapsible-fifth-title]
ldap_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_domain
|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST and PATCH.

```

```

|base_dn
|string
a|Specifies the default base DN for all searches.

|base_scope
|string
a|Specifies the default search scope for LDAP queries:

* base - search the named entry only
* onelevel - search all entries immediately below the DN
* subtree - search the named DN entry and the entire subtree below the DN

|bind_dn
|string
a|Specifies the user that binds to the LDAP servers.

|bind_password
|string
a|Specifies the bind password for the LDAP servers.

|min_bind_level
|string
a|The minimum bind authentication level. Possible values are:

* anonymous - anonymous bind
* simple - simple bind
* sasl - Simple Authentication and Security Layer (SASL) bind

|port
|integer
a|The port used to connect to the LDAP Servers.

|preferred_ad_servers
|array[string]
a|

|schema
|string
a|The name of the schema template used by the SVM.

```

- * AD-IDMU - Active Directory Identity Management for UNIX
- * AD-SFU - Active Directory Services for UNIX
- * MS-AD-BIS - Active Directory Identity Management for UNIX
- * RFC-2307 - Schema based on RFC 2307
- * Custom schema

```
|servers
|array[string]
a|
```

```
|session_security
|string
a|Specifies the level of security to be used for LDAP communications:
```

- * none - no signing or sealing
- * sign - sign LDAP traffic
- * seal - seal and sign LDAP traffic

```
|svm
|link:#svm[svm]
a|
```

```
|use_start_tls
|boolean
a|Specifies whether or not to use Start TLS over LDAP connections.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[IDd7be28826cf9185a95e753196d163764]]
```

= Delete the LDAP configuration for an SVM

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/name-services/ldap/{svm.uuid}`#
```

Introduced In: 9.6

Deletes the LDAP configuration of the specified SVM. LDAP can be removed as a source from the ns-switch if LDAP is not used as a source for lookups.

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the SVM to which this object belongs.
```

```
|===
```

== Response

Status: 200, Ok

== Error

Status: Default, Error

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string

```

```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[ID5c1986849072cb2329b242fbfd9177f9]]
= Retrieve the LDAP configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/name-services/ldap/{svm.uuid}`#
```

Introduced In: 9.6

Retrieves LDAP configuration for an SVM. All parameters for the LDAP configuration are displayed by default.

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
```

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ad_domain
|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST and PATCH.

|base_dn
|string
a|Specifies the default base DN for all searches.

|base_scope
|string
a|Specifies the default search scope for LDAP queries:

* base - search the named entry only
* onelevel - search all entries immediately below the DN
* subtree - search the named DN entry and the entire subtree below the DN

|bind_dn
|string
a|Specifies the user that binds to the LDAP servers.

|min_bind_level
|string
a|The minimum bind authentication level. Possible values are:

* anonymous - anonymous bind
* simple - simple bind
* sasl - Simple Authentication and Security Layer (SASL) bind

|port
|integer
a|The port used to connect to the LDAP Servers.

```

```

|preferred_ad_servers
|array[string]
a|

|schema
|string
a|The name of the schema template used by the SVM.

* AD-IDMU - Active Directory Identity Management for UNIX
* AD-SFU - Active Directory Services for UNIX
* MS-AD-BIS - Active Directory Identity Management for UNIX
* RFC-2307 - Schema based on RFC 2307
* Custom schema

|servers
|array[string]
a|

|session_security
|string
a|Specifies the level of security to be used for LDAP communications:

* none - no signing or sealing
* sign - sign LDAP traffic
* seal - seal and sign LDAP traffic

|svm
|link:#svm[svm]
a|

|use_start_tls
|boolean
a|Specifies whether or not to use Start TLS over LDAP connections.

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {

```

```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_domain": "string",
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "min_bind_level": "string",
  "port": "389",
  "preferred_ad_servers": [
    "string"
  ],
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]

```

```

_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

```



```

|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

|===

//end collapsible .Definitions block

====

[[IDac09cc5f1c37041437763e9722b24765]]

= Update the LDAP configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-block]#`/name-services/ldap/{svm.uuid}`#

Introduced In: 9.6

Updates an LDAP configuration of an SVM.

== Important notes

* Both mandatory and optional parameters of the LDAP configuration can be updated.

* The LDAP servers and Active Directory domain are mutually exclusive fields. These fields cannot be empty. At any point in time, either the LDAP servers or Active Directory domain must be populated.

* IPv6 must be enabled if IPv6 family addresses are specified.

Configuring more than one LDAP server is recommended to avoid a single point of failure.

Both FQDNs and IP addresses are supported for the "servers" field.

The Active Directory domain or LDAP servers are validated as part of this operation.

LDAP validation fails in the following scenarios:

- . The server does not have LDAP installed.
- . The server or Active Directory domain is invalid.
- . The server or Active Directory domain is unreachable

== Parameters

[cols=5*,options=header]

|===

|Name

```
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|ad_domain
|string
```

```
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST and PATCH.
```

```
|base_dn
|string
```

```
a|Specifies the default base DN for all searches.
```

```
|base_scope
|string
```

```
a|Specifies the default search scope for LDAP queries:
```

- * base - search the named entry only
- * onelevel - search all entries immediately below the DN
- * subtree - search the named DN entry and the entire subtree below the DN

```
|bind_dn
|string
```

```
a|Specifies the user that binds to the LDAP servers.
```

```

|bind_password
|string
a|Specifies the bind password for the LDAP servers.

|min_bind_level
|string
a|The minimum bind authentication level. Possible values are:

* anonymous - anonymous bind
* simple - simple bind
* sasl - Simple Authentication and Security Layer (SASL) bind

|port
|integer
a|The port used to connect to the LDAP Servers.

|preferred_ad_servers
|array[string]
a|

|schema
|string
a|The name of the schema template used by the SVM.

* AD-IDMU - Active Directory Identity Management for UNIX
* AD-SFU - Active Directory Services for UNIX
* MS-AD-BIS - Active Directory Identity Management for UNIX
* RFC-2307 - Schema based on RFC 2307
* Custom schema

|servers
|array[string]
a|

|session_security
|string
a|Specifies the level of security to be used for LDAP communications:

* none - no signing or sealing
* sign - sign LDAP traffic
* seal - seal and sign LDAP traffic

```

```
|svm
|link:#svm[svm]
a|

|use_start_tls
|boolean
a|Specifies whether or not to use Start TLS over LDAP connections.
```

```
|===
```

```
.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "ad_domain": "string",
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "bind_password": "string",
  "min_bind_level": "string",
  "port": "389",
  "preferred_ad_servers": [
    "string"
  ],
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

== Response
```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 262186

| LDAP Servers cannot be used with Active Directory domain and/or preferred Active Directory servers

| 2621488

| Invalid SVM context

| 2621706

| The specified SVM UUID is incorrect for the specified SVM name

| 4915203

| The specified LDAP schema does not exist

| 4915208

| The specified LDAP servers or preferred Active Directory servers contain duplicate server entries

| 4915229

| DNS resolution failed due to an internal error. Contact technical support if this issue persists

| 4915231

| DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured

| 23724132

| DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured

| 4915234

| The specified LDAP server or preferred Active Directory server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast

| 4915248

| LDAP servers cannot be empty or "-". Specified Active Directory domain

is invalid because it is empty or "-" or it contains either the special characters or "-" at the start or end of the domain.

| 4915251
| STARTTLS and LDAPS cannot be used together

| 4915257
| The LDAP configuration is invalid. Verify that the distinguished names and bind password are correct

| 4915258
| The LDAP configuration is invalid. Verify that the Active Directory domain or servers are reachable and that the network configuration is correct

| 4915259
| LDAP configurations with Active Directory domains are not supported on admin SVM.

| 23724130
| Cannot use an IPv6 name server address because there are no IPv6 LIFs
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

=====

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#ldap_service]
[.api-collapsible-fifth-title]
ldap_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_domain
|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST and PATCH.

|base_dn
|string
a|Specifies the default base DN for all searches.

```



```

|base_scope
|string
a|Specifies the default search scope for LDAP queries:

* base - search the named entry only
* onelevel - search all entries immediately below the DN
* subtree - search the named DN entry and the entire subtree below the DN


|bind_dn
|string
a|Specifies the user that binds to the LDAP servers.


|bind_password
|string
a|Specifies the bind password for the LDAP servers.


|min_bind_level
|string
a|The minimum bind authentication level. Possible values are:

* anonymous - anonymous bind
* simple - simple bind
* sasl - Simple Authentication and Security Layer (SASL) bind


|port
|integer
a|The port used to connect to the LDAP Servers.


|preferred_ad_servers
|array[string]
a|

|schema
|string
a|The name of the schema template used by the SVM.

* AD-IDMU - Active Directory Identity Management for UNIX
* AD-SFU - Active Directory Services for UNIX
* MS-AD-BIS - Active Directory Identity Management for UNIX
* RFC-2307 - Schema based on RFC 2307

```

* Custom schema

|servers

|array[string]

a|

|session_security

|string

a|Specifies the level of security to be used for LDAP communications:

* none - no signing or sealing

* sign - sign LDAP traffic

* seal - seal and sign LDAP traffic

|svm

|link:#svm[svm]

a|

|use_start_tls

|boolean

a|Specifies whether or not to use Start TLS over LDAP connections.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage name mappings for SVMs

```

```
:leveloffset: +1
```

```
[[IDdd25b9cb894204cf153f6938eb160dfc]]  
= Name-services name-mappings endpoint overview
```

== Overview

Name mapping is used to map CIFS identities to UNIX identities, Kerberos identities to UNIX identities, and UNIX identities to CIFS identities. It needs this information to obtain user credentials and provide proper file access regardless of whether they are connecting from an NFS client or a CIFS client.

The system keeps a set of conversion rules for each Storage Virtual Machine (SVM). Each rule consists of two pieces: a pattern and a replacement. Conversions start at the beginning of the appropriate list and perform a substitution based on the first matching rule. The pattern is a UNIX-style regular expression. The replacement is a string containing escape sequences representing subexpressions from the pattern, as in the UNIX sed program.

Name mappings are applied in the order in which they occur in the priority list; for example, a name mapping that occurs at position 2 in the priority list is applied before a name mapping that occurs at position 3. Each mapping direction (Kerberos-to-UNIX, Windows-to-UNIX, and UNIX-to-Windows) has its own priority list. You are prevented from creating two name mappings with the same pattern.

== Examples

=== Creating a name-mapping with client_match as the ip-address

Use the following API to create a name-mapping. Note the `_return_records=true_` query parameter is used to obtain the newly created entry in the response.

The API:

```
POST /api//name-services/name-mappings
```

The call:

```
curl -X POST "https://<mgmt-ip>/api/name-services/name-  
mappings?return_records=true" -H "accept: application/json" -H "Content-
```

```
Type: application/json" -d "{ \"client_match\": \"10.254.101.111/28\",
\"direction\": \"win_unix\", \"index\": 1, \"pattern\":
\"ENGCIIFS_AD_USER\", \"replacement\": \"unix_user1\", \"svm\": { \"name\":
\"vs1\", \"uuid\": \"f71d3640-0226-11e9-8526-000c290a8c4b\" }}"
```

The response:

```
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
        "name": "vs1"
      },
      "direction": "win_unix",
      "index": 1,
      "pattern": "ENGCIIFS_AD_USER",
      "replacement": "unix_user1",
      "client_match": "10.254.101.111/28"
    }
  ]
}
```

=== Creating a name-mapping with client_match as the hostname

Use the following API to create a name-mapping. Note the `_return_records=true_` query parameter is used to obtain the newly created entry in the response.

The API:

POST /api//name-services/name-mappings

The call:

```
curl -X POST "https://<mgmt-ip>/api/name-services/name-
mappings?return_records=true" -H "accept: application/json" -H "Content-
Type: applicatio/json" -d "{ \"client_match\": \"google.com\",
\"direction\": \"win_unix\", \"index\": 2, \"pattern\":
\"ENGCIIFS_AD_USER\", \"replacement\": \"unix_user1\", \"svm\": { \"name\":
\"vs1\", \"uuid\": \"f71d3640-0226-11e9-8526-000c290a8c4b\" }}"
```

The response:

```
{
  "num_records": 1,
```

```

"records": [
  {
    "svm": {
      "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
      "name": "vs1"
    },
    "direction": "win_unix",
    "index": 2,
    "pattern": "ENGCIIFS_AD_USER",
    "replacement": "unix_user1",
    "client_match": "google.com"
  }
]
}
----

```

=== Retrieving all name-mapping configurations for all SVMs in the cluster

The API:

GET /api/name-services/name-mappings

The call:

```

curl -X GET "https://<mgmt-ip>/api/name-services/name-
mappings?fields=&return_records=true&return_timeout=15" -H "accept:
application/json"

```

The response:

```

{
  "records": [
    {
      "svm": {
        "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
        "name": "vs1"
      },
      "direction": "win_unix",
      "index": 1,
      "pattern": "ENGCIIFS_AD_USER",
      "replacement": "unix_user1",
      "client_match": "10.254.101.111/28"
    },
    {
      "svm": {
        "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
        "name": "vs1"
      }
    }
  ]
}

```

```

    },
    "direction": "win_unix",
    "index": 2,
    "pattern": "ENGCIIFS_AD_USER",
    "replacement": "unix_user1",
    "client_match": "google.com"
  }
],
"num_records": 2
}
----
```

=== Retrieving a name-mapping configuration for a specific SVM, and for the specified direction and index

```
'''
```

```
----
```

The API:

```
GET /api/name-services/name-mappings/{svm.uuid}/{direction}/{index}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/name-services/name-mappings/f71d3640-0226-11e9-8526-000c290a8c4b/win_unix/1" -H "accept: application/json"
```

The response:

```

{
  "svm": {
    "uuid": "f71d3640-0226-11e9-8526-000c290a8c4b",
    "name": "vs1"
  },
  "direction": "win_unix",
  "index": 1,
  "pattern": "ENGCIIFS_AD_USER",
  "replacement": "unix_user1",
  "client_match": "10.254.101.111/28"
}
----
```

```
'''
```

=== Updating a specific name-mapping configuration

```
'''
```

```
----
```

```

# The API:
PATCH /api//name-services/name-mappings/{svm.uuid}/{direction}/{index}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/name-mappings/f71d3640-0226-11e9-8526-000c290a8c4b/win_unix/1" -H "accept: application/json" -H "Content-Type: application/json" -d '{"client_match": "\10.254.101.222/28", "pattern": "ENGCIIFS_LOCAL_USER", "replacement": "pcuser"}'

# swapping a specified namemapping entry by index
curl -X PATCH "https://<mgmt-ip>/api/name-services/name-mappings/f71d3640-0226-11e9-8526-000c290a8c4b/win-unix/3?new_index=1" -H "accept: application/json" -H "Content-Type: application/json" -d '{"pattern": "ENGCIIFS_AD_USER", "replacement": "unix_user1"}'
----

'''

=== Removing a specific name-mapping configuration

'''

----

# The API:
DELETE /api/name-services/name-mappings/{svm.uuid}/{direction}/{index}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/name-mappings/f71d3640-0226-11e9-8526-000c290a8c4b/win_unix/1" -H "accept: application/json"
----

'''

[[ID599bf4b91152a03e252e926cb1ef3d0a]]
= Retrieve the hostname mapping for all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/name-services/name-mappings`#

```


Introduced In: 9.6

Retrieves the name mapping configuration for all SVMs.

== Related ONTAP commands

* `vserver name-mapping show`

== Learn more

* xref:{relative_path}name-services_name-mappings_endpoint_overview.html[DOC /name-services/name-mappings]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|direction

|string

|query

|False

a|Filter by direction

|replacement

|string

|query

|False

a|Filter by replacement

|index

|integer

|query

|False

a|Filter by index

|pattern

```
|string
|query
|False
a|Filter by pattern
```

```
|client_match
|string
|query
|False
a|Filter by client_match
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

|return_timeout
|integer
|query
|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

```

|records
|array[link:#name_mapping[name_mapping]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "client_match": "10.254.101.111/28",
      "direction": "win_unix",
      "index": "1",
      "pattern": "ENGCIIFS_AD_USER",
      "replacement": "unix_user1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
=====

```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]

```

```

a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#name_mapping]
[.api-collapsible-fifth-title]
name_mapping

Name mapping is used to map CIFS identities to UNIX identities, Kerberos
identities to UNIX identities, and UNIX identities to CIFS identities. It
needs this information to obtain user credentials and provide proper file
access regardless of whether they are connecting from an NFS client or a
CIFS client.

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|_links
|link:#_links[_links]
a|

|client_match
|string
a|Client workstation IP Address which is matched when searching for the
pattern.
You can specify the value in any of the following formats:

* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
* As a hostname

|direction
|string
a|Direction in which the name mapping is applied. The possible values are:

*** krb_unix - Kerberos principal name to UNIX user name

*** win_unix - Windows user name to UNIX user name

*** unix_win - UNIX user name to Windows user name mapping

|index
|integer
a|Position in the list of name mappings.

|pattern
|string
a|Pattern used to match the name while searching for a name that can be
used as a replacement. The pattern is a UNIX-style regular expression.
Regular expressions are case-insensitive when mapping from Windows to
UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and
UNIX to Windows.

|replacement

```



```

|string
a|The name that is used as a replacement, if the pattern associated with
this entry matches.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]

```

```
a|Message arguments
```

```
|code  
|string  
a|Error code
```

```
|message  
|string  
a|Error message
```

```
|target  
|string  
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block  
====
```

```
[[IDf7f25a37086d887b9542bbe7fab03a0e]]  
= Create hostname mappings for an SVM
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-  
block]#`/name-services/name-mappings`#
```

Introduced In: 9.6

Creates name mappings for an SVM.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the name mapping.

* `index` - Name mapping's position in the priority list.

* `direction` - Direction of the name mapping.

* `pattern` - Pattern to match to. Maximum length is 256 characters.

* `replacement` - Replacement pattern to match to. Maximum length is 256 characters.

== Recommended optional properties

* `client_match` - Hostname or IP address added to match the pattern to the client's workstation IP address.

== Related ONTAP commands

* `vserver name-mapping create`
* `vserver name-mapping insert`

== Learn more

* xref:{relative_path}name-services_name-mappings_endpoint_overview.html[DOC /name-services/name-mappings]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|client_match

```

|string
a|Client workstation IP Address which is matched when searching for the
pattern.
You can specify the value in any of the following formats:

* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
* As a hostname

|direction
|string
a|Direction in which the name mapping is applied. The possible values are:

*** krb_unix - Kerberos principal name to UNIX user name

*** win_unix - Windows user name to UNIX user name

*** unix_win - UNIX user name to Windows user name mapping

|index
|integer
a|Position in the list of name mappings.

|pattern
|string
a|Pattern used to match the name while searching for a name that can be
used as a replacement. The pattern is a UNIX-style regular expression.
Regular expressions are case-insensitive when mapping from Windows to
UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and
UNIX to Windows.

|replacement
|string
a|The name that is used as a replacement, if the pattern associated with
this entry matches.

|svm
|link:#svm[svm]

```

a|

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "client_match": "10.254.101.111/28",
  "direction": "win_unix",
  "index": "1",
  "pattern": "ENGCIIFS_AD_USER",
  "replacement": "unix_user1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

====

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#name_mapping[name_mapping]]
a|

|===


.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "client_match": "10.254.101.111/28",
      "direction": "win_unix",
      "index": "1",
      "pattern": "ENGCIIFS_AD_USER",
      "replacement": "unix_user1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 65798185
| Failed to resolve the specified hostname

| 65798149
| Invalid index for the name mapping entry

| 2621706
| The specified svm.uuid and svm.name refer to different SVMs
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

```

```

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

```

|===

```

```

[#name_mapping]
[.api-collapsible-fifth-title]
name_mapping

```

Name mapping is used to map CIFS identities to UNIX identities, Kerberos identities to UNIX identities, and UNIX identities to CIFS identities. It needs this information to obtain user credentials and provide proper file access regardless of whether they are connecting from an NFS client or a CIFS client.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|client_match
|string
a|Client workstation IP Address which is matched when searching for the
pattern.

```

You can specify the value in any of the following formats:

- * As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24
- * As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8ble:b255:4071::/64
- * As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0
- * As a hostname


```

|direction
|string
a|Direction in which the name mapping is applied. The possible values are:

*** krb_unix - Kerberos principal name to UNIX user name

*** win_unix - Windows user name to UNIX user name

*** unix_win - UNIX user name to Windows user name mapping


|index
|integer
a|Position in the list of name mappings.


|pattern
|string
a|Pattern used to match the name while searching for a name that can be
used as a replacement. The pattern is a UNIX-style regular expression.
Regular expressions are case-insensitive when mapping from Windows to
UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and
UNIX to Windows.


|replacement
|string
a|The name that is used as a replacement, if the pattern associated with
this entry matches.


|svm
|link:#svm[svm]
a|

|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments


[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

```

```
//end collapsible .Definitions block
====

[[ID54b977788c5080177664f9bcd0e323f2]]
= Delete the name mapping configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/name-services/name-mappings/{svm.uuid}/{direction}/{index}`#

*Introduced In:* 9.6

Deletes the name mapping configuration.

== Related ONTAP commands

* `vserver name-mapping delete`

== Learn more

* xref:{relative_path}name-services_name-
mappings_endpoint_overview.html[DOC /name-services/name-mappings]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|direction
|string
|path
|True
a|Direction

|index
|integer
```

```
|path
|True
a|Position of the entry in the list

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ]
  }
}
```

```

    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

```

```

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID1e9b859cc0dce9c49c52ec88cb7e59b7]]
= Retrieve the name mapping configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/name-services/name-mappings/{svm.uuid}/{direction}/{index}`#
```

Introduced In: 9.6

Retrieves the name mapping configuration of an SVM.

== Related ONTAP commands

* `vserver name-mapping show`

== Learn more

* xref:{relative_path}name-services_name-
mappings_endpoint_overview.html[DOC /name-services/name-mappings]

```
== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|direction
|string
|path
|True
a|Direction

|index
|integer
|path
|True
a|Position of the entry in the list

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|client_match
|string
a|Client workstation IP Address which is matched when searching for the
pattern.
You can specify the value in any of the following formats:

* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
* As a hostname

|direction
|string
a|Direction in which the name mapping is applied. The possible values are:

*** krb_unix - Kerberos principal name to UNIX user name

*** win_unix - Windows user name to UNIX user name

*** unix_win - UNIX user name to Windows user name mapping

|index
|integer
a|Position in the list of name mappings.

|pattern
|string
a|Pattern used to match the name while searching for a name that can be
used as a replacement. The pattern is a UNIX-style regular expression.
Regular expressions are case-insensitive when mapping from Windows to
UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and
```


UNIX to Windows.

```
|replacement
|string
a|The name that is used as a replacement, if the pattern associated with
this entry matches.
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "client_match": "10.254.101.111/28",
  "direction": "win_unix",
  "index": "1",
  "pattern": "ENGCIFS_AD_USER",
  "replacement": "unix_user1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

====

== Error

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid

```

```
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

=====

[[IDb606d0667d2222b3fe74e41aa85fcb1a]]

= Update the name mapping configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-block]#`/name-services/name-mappings/{svm.uuid}/{direction}/{index}`#

Introduced In: 9.6

Updates the name mapping configuration of an SVM. The positions can be swapped by providing the ``new_index`` property.

Swapping is not allowed for entries that have ``client_match`` property configured.

== Related ONTAP commands

* ``vserver name-mapping insert``

* ``vserver name-mapping modify``

* ``vserver name-mapping swap``

== Learn more

* xref:{relative_path}name-services_name-mappings_endpoint_overview.html[DOC /name-services/name-mappings]

== Parameters

```

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|direction
|string
|path
|True
a|Direction

|index
|integer
|path
|True
a|Position of the entry in the list

|new_index
|integer
|query
|False
a|New position of the Index after a swap is completed.

* Introduced in: 9.7

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|client_match
|string
a|Client workstation IP Address which is matched when searching for the
pattern.
You can specify the value in any of the following formats:

* As an IPv4 address with a subnet mask expressed as a number of bits; for
instance, 10.1.12.0/24
* As an IPv6 address with a subnet mask expressed as a number of bits; for
instance, fd20:8ble:b255:4071::/64
* As an IPv4 address with a network mask; for instance,
10.1.16.0/255.255.255.0
* As a hostname

|pattern
|string
a|Pattern used to match the name while searching for a name that can be
used as a replacement. The pattern is a UNIX-style regular expression.
Regular expressions are case-insensitive when mapping from Windows to
UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and
UNIX to Windows.

|replacement
|string
a|The name that is used as a replacement, if the pattern associated with
this entry matches.

|svm
|link:#svm[svm]
a|

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "client_match": "10.254.101.111/28",

```

```

"pattern": "ENGCIIFS_AD_USER",
"replacement": "unix_user1",
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 65798185
| Failed to resolve the specified hostname
|
| 65798179
| Cannot swap entries because one or both entries have hostname or address
| configured.
|
| Delete and recreate the new entry at the specified position.
|====

```

```

schema:
$ref: "#/definitions/error_response"

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]

```



```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the SVM.
```

```
|===
```

```
[#name_mapping]
```

```
[.api-collapsible-fifth-title]
```

```
name_mapping
```

Name mapping is used to map CIFS identities to UNIX identities, Kerberos

identities to UNIX identities, and UNIX identities to CIFS identities. It needs this information to obtain user credentials and provide proper file access regardless of whether they are connecting from an NFS client or a CIFS client.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|client_match
```

```
|string
```

a|Client workstation IP Address which is matched when searching for the pattern.

You can specify the value in any of the following formats:

- * As an IPv4 address with a subnet mask expressed as a number of bits; for instance, 10.1.12.0/24

- * As an IPv6 address with a subnet mask expressed as a number of bits; for instance, fd20:8ble:b255:4071::/64

- * As an IPv4 address with a network mask; for instance, 10.1.16.0/255.255.255.0

- * As a hostname

```
|pattern
```

```
|string
```

a|Pattern used to match the name while searching for a name that can be used as a replacement. The pattern is a UNIX-style regular expression. Regular expressions are case-insensitive when mapping from Windows to UNIX, and they are case-sensitive for mappings from Kerberos to UNIX and UNIX to Windows.

```
|replacement
```

```
|string
```

a|The name that is used as a replacement, if the pattern associated with this entry matches.

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
= Manage NIS server configurations
```

```
:leveloffset: +1
```

```
[[IDb616f4f43dce8ad10e11e672715fa81f]]
= Name-services NIS endpoint overview
```

```
== Overview
```

NIS servers are used to authenticate user and client computers. NIS domain name and NIS server information is required to configure NIS. It is important to note that this API is used to retrieve and manage NIS server configurations for data SVMs only. NIS configuration for the cluster is managed via `xref:{relative_path}security-authentication-cluster-nis-endpoint-overview.html[/api/security/authentication/cluster/nis]` .

```
== Retrieving NIS Information
```

The NIS GET endpoint retrieves all of the NIS configurations for data SVMs.

```
== Examples
```

```
=== Retrieving all fields for all NIS configurations
```

```
'''
```

```
----
```

```
# The API:
/api/name-services/nis
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/name-services/nis?fields=*" -H "accept: application/hal+json"
```

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
          }
        }
      },
      "domain": "domainA.example.com",
      "servers": [
        "10.10.10.10",
        "example.com"
      ],
      "bound-servers": [
        "10.10.10.10"
      ],
      "_links": {
        "self": {
          "href": "/api/name-services/nis/179d3c85-7053-11e8-b9b8-005056b41bd1"
        }
      }
    },
    {
      "svm": {
        "uuid": "6a52023b-7066-11e8-b9b8-005056b41bd1",
        "name": "vs2",
        "_links": {
          "self": {
            "href": "/api/svm/svms/6a52023b-7066-11e8-b9b8-005056b41bd1"
          }
        }
      },
      "domain": "domainB.example.com",
      "servers": [
        "2.2.2.2",
        "3.3.3.3",
        "4.4.4.4"
      ]
    }
  ]
}
```

```

    ],
    "bound-servers": [],
    "_links": {
      "self": {
        "href": "/api/name-services/nis/6a52023b-7066-11e8-b9b8-005056b41bd1"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/name-services/nis?fields=*"
  }
}
}
}
----

'''

=== Retrieving all NIS configurations whose bound servers start with _10_

'''

----

# The API:
/api/name-services/nis

# The call:
curl -X GET "https://<mgmt-ip/api/name-services/nis?bound_servers=10*" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/179d3c85-7053-11e8-b9b8-005056b41bd1"
          }
        }
      }
    }
  ]
}

```

```

    },
    "bound-servers": [
        "10.10.10.10"
    ],
    "_links": {
        "self": {
            "href": "/api/name-services/nis/6a52023b-7066-11e8-b9b8-005056b41bd1"
        }
    }
},
"num_records": 1,
"_links": {
    "self": {
        "href": "/api/name-services/nis?bound_servers=10*"
    }
}
}
}

```

'''

=== Retrieving the NIS configuration of a specific SVM

'''

The API:

/api/name-services/nis/{svm.uuid}

The call:

```
curl -X GET "https://<mgmt-ip>/api/name-services/nis/179d3c85-7053-11e8-b9b8-005056b41bd1" -H "accept: application/hal+json"
```

The response:

```

{
  "svm": {
    "uuid": "179d3c85-7053-11e8-b9b8-005056b41bd1",
    "name": "vs1"
  },
  "domain": "domainA.example.com",
  "servers": [
    "10.10.10.10",
    "example.com"
  ]
}

```

```

],
"bound_servers": [
    "10.10.10.10"
]
}
-----

```

'''

== Creating a NIS configuration

The NIS POST endpoint creates a NIS configuration for the specified SVM.

== Example

The following example shows a POST operation:

The API:

/api/name-services/nis

The call:

```

curl -X POST "https://<mgmt-ip>/api/name-services/nis" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"svm\": {
\"uuid\": \"179d3c85-7053-11e8-b9b8-005056b41bd1\" }, \"domain\":
\"domainA.example.com\", \"servers\": [ \"10.10.10.10\", \"example.com\"
] }"
-----

```

'''

== Updating the NIS configuration

The NIS PATCH endpoint updates the NIS configuration for the specified NIS server.

== Examples

=== Updating the domain

'''

The API:

/api/name-services/nis/{svm.uuid}

```
# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/nis/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"domain\": \"domainC.example.com\", \"servers\":
[ \"13.13.13.13\" ]}"
```

'''

== Updating the server

'''

```
# The API:
/api/name-services/nis/{svm.uuid}
```

```
# The call:
curl -X PATCH "https://<mgmt-ip>/api/name-services/nis/179d3c85-7053-11e8-
b9b8-005056b41bd1" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"servers\": [ \"14.14.14.14\" ]}"
```

'''

== Deleting a NIS configuration

The NIS DELETE endpoint deletes the NIS configuration for the specified SVM.

== Example

The following example shows a DELETE operation:

'''

```
# The API:
/api/name-services/nis/{svm.uuid}
```

```
# The call:
curl -X DELETE "https://<mgmt-ip>/api/name-services/nis/179d3c85-7053-
11e8-b9b8-005056b41bd1" -H "accept: application/hal+json"
```

'''

[[ID4809bb0037d246ce278e11db435dea6d]]

= Retrieve NIS domain configurations of all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/name-services/nis`#

Introduced In: 9.6

Retrieves NIS domain configurations of all the SVMs. The `bound_servers` field indicates the successfully bound NIS servers. Lookups and authentications fail if there are no bound servers.

== Related ONTAP commands

* ``vserver services name-service nis-domain show``

* ``vserver services name-service nis-domain show-bound``

== Learn more

* `xref:{relative_path}name-services_nis_endpoint_overview.html` [DOC `/name-services/nis`]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|domain

|string

|query

|False

a|Filter by domain

* Introduced in: 9.7

|servers
|string
|query
|False
a|Filter by servers

* Introduced in: 9.7

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

* Introduced in: 9.7

|svm.name
|string
|query
|False
a|Filter by svm.name

* Introduced in: 9.7

|bound_servers
|string
|query
|False
a|Filter by bound_servers

* Introduced in: 9.7

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records

```

|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of NIS domain records.

|records
|array[link:#nis_service[nis_service]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "bound_servers": [
        "string"
      ],
      "domain": "string",
      "servers": [
        "string"
      ],
    },
  ],
}

```

```

    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",

```

```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

```

```

[#nis_service]
[.api-collapsible-fifth-title]
nis_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|bound_servers
|array[string]
a|

|domain
|string
a|The NIS domain to which this configuration belongs.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code

```



```

|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

[[ID3a0577ad60259ca90a2e7b14ff7920ec]]

= Create the NIS domain and server configuration for a data SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/name-services/nis`#

Introduced In: 9.6

Creates an NIS domain and server configuration for a data SVM.

NIS configuration for the cluster is managed via

xref:{relative_path}security-authentication-cluster-nis-endpoint-overview.html[/api/security/authentication/cluster/nis] .

== Important notes

- * Each SVM can have one NIS domain configuration.

- * Multiple SVMs can be configured with the same NIS domain. Specify the NIS domain and NIS servers as input. Domain name and servers fields cannot be empty.

- * Both FQDNs and IP addresses are supported for the servers field.

- * IPv6 must be enabled if IPv6 family addresses are specified in the servers field.

- * A maximum of ten NIS servers are supported.

== Required properties

- * `svm.uuid` or `svm.name` - Existing SVM in which to create the NIS configuration.

- * `domain` - NIS domain to which the configuration belongs.

- * `servers` - List of NIS server IP addresses.

== Related ONTAP commands

- * `vserver services name-service nis-domain create`

== Learn more

- * xref:{relative_path}name-services_nis_endpoint_overview.html[DOC /name-services/nis]

== Parameters

[cols=5*,options=header]

```

|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|bound_servers
|array[string]
a|

|domain
|string
a|The NIS domain to which this configuration belongs.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|svm
|link:#svm[svm]
a|

```

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "bound_servers": [
    "string"
  ],
  "domain": "string",
  "servers": [
    "string"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

====

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of NIS domain records.

|records
|array[link:#nis_service[nis_service]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "bound_servers": [
        "string"
      ],
      "domain": "string",
      "servers": [
        "string"
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 1966253

| IPv6 is not enabled in the cluster

| 2621488

| Invalid SVM context

| 2621706

| The specified SVM UUID is incorrect for the specified SVM name

| 3276964

| NIS domain name or NIS server domain is too long. The maximum supported for domain name is 64 characters and the maximum supported for NIS server domain is 255 characters

| 3276933

| A maximum of 10 NIS servers can be configured per SVM

| 13434916

| The SVM is in the process of being created. Wait a few minutes, and then try the command again.

| 23724109

| DNS resolution failed for one or more specified servers

| 23724112

| DNS resolution failed due to an internal error. Contact technical support if this issue persists

| 23724132

| DNS resolution failed for all the specified servers

| 23724130

| Cannot use an IPv6 name server address because there are no IPv6 LIFs

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#nis_service]
[.api-collapsible-fifth-title]
nis_service

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|bound_servers
|array[string]
a|

|domain
|string
a|The NIS domain to which this configuration belongs.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```



```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID50530db54e72b1f874e91d63f403f1ec]]
= Delete the NIS domain configuration for an SVM

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
```

```
block]#`/name-services/nis/{svm.uuid}`#
```

Introduced In: 9.6

Deletes the NIS domain configuration of an SVM. NIS can be removed as a source from ns-switch if NIS is not used for lookups.

== Related ONTAP commands

* `vserver services name-service nis-domain delete`

== Learn more

* xref:{relative_path}name-services_nis_endpoint_overview.html[DOC /name-services/nis]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|===
```

== Response

Status: 200, Ok

== Error

Status: Default, Error

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

```
//end collapsible .Definitions block
```

```
====
```

```
[[IDbbf4bcf797100ce77b19554fda785e58]]
```

```
= Retrieve the NIS domain and server configurations for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/name-services/nis/{svm.uuid}`#
```

```
*Introduced In:* 9.6
```

Retrieves NIS domain and server configurations of an SVM. Both NIS domain and servers are displayed by default. The `bound_servers` field indicates the successfully bound NIS servers.

```
== Related ONTAP commands
```

```
* `vserver services name-service nis-domain show`
```

```
* `vserver services name-service nis-domain show-bound`
```

```
== Learn more
```

```
* xref:{relative_path}name-services_nis_endpoint_overview.html[DOC /name-  
services/nis]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|==
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|bound_servers
|array[string]
a|

|domain
|string
a|The NIS domain to which this configuration belongs.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
```

```

=====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bound_servers": [
    "string"
  ],
  "domain": "string",
  "servers": [
    "string"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error

```

```

[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```



```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

```

```
//end collapsible .Definitions block
====
```

```
[[ID299d9b5fdbb3dc17495d6b6892c13ba1]]
```

= Update the NIS domain and server configuration for an SVM

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-  
block]#`/name-services/nis/{svm.uuid}`#
```

Introduced In: 9.6

Updates NIS domain and server configuration of an SVM.

== Important notes

- * Both NIS domain and servers can be modified.
- * Domains and servers cannot be empty.
- * Both FQDNs and IP addresses are supported for the servers field.
- * If the domain is modified, NIS servers must also be specified.
- * IPv6 must be enabled if IPv6 family addresses are specified for the servers field.

== Related ONTAP commands

* `vserver services name-service nis-domain modify`

== Learn more

* xref:{relative_path}name-services_nis_endpoint_overview.html[DOC /name-services/nis]

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|bound_servers
|array[string]
a|

|domain
|string
a|The NIS domain to which this configuration belongs.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|svm
|link:#svm[svm]
a|

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "bound_servers": [
    "string"

```

```

],
"domain": "string",
"servers": [
  "string"
],
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
=====

```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 1966253
| IPv6 is not enabled in the cluster

| 2621488
| Invalid SVM context

| 2621706
| The specified SVM UUID is incorrect for the specified SVM name

| 3276964
| NIS domain name or NIS server domain is too long. The maximum supported
for domain name is 64 characters and the maximum supported for NIS server
domain is 255 characters

| 3276933
| A maximum of 10 NIS servers can be configured per SVM

| 23724109
| DNS resolution failed for one or more specified servers

```

```
| 23724112
| DNS resolution failed due to an internal error. Contact technical
support if this issue persists

| 23724132
| DNS resolution failed for all the specified servers

| 23724130
| Cannot use an IPv6 name server address because there are no IPv6 LIFs
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
```

```

|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#nis_service]
[.api-collapsible-fifth-title]
nis_service

[cols=3*,options=header]
|===
|Name
|Type
|Description

|bound_servers
|array[string]
a|

|domain
|string
a|The NIS domain to which this configuration belongs.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|svm
|link:#svm[svm]
a|

|===

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

```



```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
:leveloffset: -1
```

```
:leveloffset: -1
```

```
= Networking
```

```
:leveloffset: +1
```

```
[[IDce4ebfa0af68b27d30cf2bbc732dfb9d]]
= Networking overview
```

```
== Overview
```

The ONTAP networking APIs enable reporting on networking information, such as IPspaces, interfaces, routes, ports, service policies and broadcast domains. Some can also be used to manage networking.

```
'''
```

```
== IPspaces
```

IPspaces enable you to configure a single ONTAP cluster so that it can be accessed by clients from more than one administratively separate network domain, even if those clients are using the same IP address subnet range. This allows for separation of client traffic for privacy and security.

An IPspace defines a distinct IP address space in which storage virtual machines (SVMs) reside. Ports and IP addresses defined for an IPspace are

applicable only within that IPspace. A distinct routing table is maintained for each SVM within an IPspace, so that no cross-SVM or cross-IPspace traffic routing occurs.

'''

== Ethernet

=== Broadcast Domains

A broadcast domain is a set of ports which would all receive a broadcast packet that is sent from any of the ports.

By accurately representing the physical network with an ONTAP broadcast domain, ONTAP ensures that IP interfaces are able to migrate to appropriate ports in case of failure. ONTAP also ensures that characteristics, such as MTU, stay matched across all ports of the broadcast domain.

A broadcast domain resides in an IPspace, and can be used by cluster-scoped or SVM-scoped IP interfaces in that IPspace. The scope of the broadcast domain's uniqueness is the IPspace it is in. You must create as many broadcast domains in an IPspace as there are IP subnets with interfaces in that IPspace.

Ports are mapped to an IPspace by assigning the port's broadcast domain.

=== Ports

A port is a physical or virtual Ethernet network device. Physical ports may be combined into Link Aggregation Groups (LAGs, or ifgrps), or divided into Virtual LANS (VLANs).

The GET and PATCH APIs are available for all port types. The POST and DELETE APIs are available for "lag" and "vlan" port types.

A given port can host zero or more IP interfaces.

A port exists in a broadcast domain and all ports within the same broadcast domain must have layer 2 network connectivity to one another. If a port within a broadcast domain goes down, any IP interfaces hosted by that port can fail over to other ports in the same broadcast domain.

'''

== Fibre Channel

=== Interfaces

Fibre Channel (FC) interfaces are the logical endpoints for Fibre Channel network connections to an SVM. A Fibre Channel interface provides Fibre Channel access to storage within the interface's SVM using either Fibre Channel Protocol (FCP) or Non-Volatile Memory Express over Fibre Channel (NVMe over FC).

The Fibre Channel interface REST API allows you to create, delete, update and discover Fibre Channel interfaces and obtain status information for Fibre Channel interfaces.

A Fibre Channel interface is created on a Fibre Channel port that is located on a cluster node. The Fibre Channel port must be specified to identify the location of the interface for a POST or PATCH that relocates an interface. You can identify the port by either supplying the node and port names or the port UUID.

=== Ports

Fibre Channel ports are the physical ports of Fibre Channel adapters on ONTAP cluster nodes that can be connected to Fibre Channel networks to provide Fibre Channel network connectivity. A Fibre Channel port defines the location of a Fibre Channel interface within the ONTAP cluster.

The Fibre Channel port REST API allows you to discover Fibre Channel ports, obtain status information for Fibre Channel ports, and configure Fibre Channel port properties.

=== Learn More

* `_Fibre Channel Logins_` found in both the `_SAN_` and `_NVMe_` sections. Fibre Channel logins represent connections formed by Fibre Channel initiators that have successfully logged in to ONTAP.

'''

== IP

=== Interfaces

An interface (also referred to as a `_LIF_` in ONTAP documentation) represents a network access point to a node in a cluster. In other words, an interface is essentially an IPv4 or IPv6 address with associated attributes.

IP interfaces are configured on ports to send and receive communications

over the network. The port that will host the interface can either be explicitly specified using node and/or port fields or implicitly specified using a broadcast domain.

The IPspace of the interface is required for cluster-scoped interfaces. The SVM owning the interface is required for SVM-scoped interfaces. These interfaces are in the SVM's IPspace.

The service policy of an interface defines what network services are provided by the interface.

=== Routes

Routes indicate which IPv4 or IPv6 gateway to use to communicate with hosts that are not on the local subnet. Typically, an IP interface (or LIF) can only use a gateway if it has the same address family and is in the LIF's subnet.

It is important that every gateway address belongs to a physical or virtual router that has connectivity to the specified destination network.

SVM-scoped routes can only be used by IP interfaces of the specified SVM. Likewise, cluster-scoped routes can only be used by cluster-scoped IP interfaces in the specified IPspace.

=== Service Policies

Service policies are named groupings that define what services are supported by an IP interface. These include both built-in service policies (for example: default-data-files or default-management) and custom service policies.

Service policies are scoped to either an SVM or IPspace.

'''

= Manage broadcast domains

:leveloffset: +1

[[IDcead08b456e4cb5b6cd15c80e2ac4f8d]]

= Network Ethernet broadcast-domains endpoint overview

== Overview

A broadcast domain is a collection of Ethernet ports that have layer 2 connectivity. They are used to determine which Ethernet ports can host interfaces of various types. The broadcast domain REST API allows you to retrieve, create, modify, and delete broadcast domains. The broadcast domain APIs do not manage port membership. To add a port to a broadcast domain or to move a port to a different broadcast domain, use PATCH `/network/ethernet/ports/+++<uuid>+++.``+++</uuid>+++`

== Retrieving network Ethernet broadcast domain information

The broadcast domains GET API retrieves and displays relevant information pertaining to the broadcast domains configured in the cluster. The API retrieves the list of all broadcast domains configured in the cluster, or a specific broadcast domain.

'''

== Examples

=== Retrieving all broadcast domains in the cluster

The following output shows the list of all broadcast domains configured in a cluster.

'''

The API:

`/api/network/ethernet/broadcast-domains`

The call:

```
curl -X GET "https://<mgmt-ip>/api/network/ethernet/broadcast-domains" -H
"accept: application/hal+json"
```

The response:

```
{
  "records": [
    {
      "uuid": "6970c2a9-f34f-11e8-8373-005056bb6b85",
      "name": "Cluster",
      "ipspace": {
```

```

    "uuid": "6267eff8-f34f-11e8-8373-005056bb6b85",
    "name": "Cluster",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/6267eff8-f34f-11e8-8373-005056bb6b85"
      }
    },
    "ports": [
      {
        "uuid": "626b4d19-f34f-11e8-8373-005056bb6b85",
        "name": "e0a",
        "node": {
          "name": "examplecluster-node01"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/626b4d19-f34f-11e8-8373-005056bb6b85"
          }
        }
      },
      {
        "uuid": "626b77b9-f34f-11e8-8373-005056bb6b85",
        "name": "e0b",
        "node": {
          "name": "examplecluster-node01"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/626b77b9-f34f-11e8-8373-005056bb6b85"
          }
        }
      }
    ],
    "mtu": 9000,
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/6970c2a9-f34f-11e8-8373-005056bb6b85"
      }
    }
  },
  {

```

```

"uuid": "6972416c-f34f-11e8-8373-005056bb6b85",
"name": "Default",
"ipSPACE": {
  "uuid": "5f650349-f34f-11e8-8373-005056bb6b85",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipSpaces/5f650349-f34f-11e8-8373-
005056bb6b85"
    }
  }
},
"ports": [
  {
    "uuid": "626bae19-f34f-11e8-8373-005056bb6b85",
    "name": "e0c",
    "node": {
      "name": "examplecluster-node01"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/626bae19-f34f-11e8-8373-
005056bb6b85"
      }
    }
  },
  {
    "uuid": "626bd677-f34f-11e8-8373-005056bb6b85",
    "name": "e0d",
    "node": {
      "name": "examplecluster-node01"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/626bd677-f34f-11e8-8373-
005056bb6b85"
      }
    }
  }
],
"mtu": 1500,
"_links": {
  "self": {
    "href": "/api/network/ethernet/broadcast-domains/6972416c-f34f-
11e8-8373-005056bb6b85"
  }
}

```

```

    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/network/ethernet/broadcast-domains?fields=*"
  }
}
}
}
----

'''

```

=== Retrieving a specific broadcast domain

The following output shows the response returned when a specific broadcast domain is requested. The system returns an error if there is no broadcast domain with the requested UUID.

```
'''
```

```
----
```

The API:

```
/api/network/ethernet/broadcast-domains/{uuid}
```

The call:

```
curl -X GET "https://<mgmt-ip>/api/network/ethernet/broadcast-
domains/4475a2c8-f8a0-11e8-8d33-005056bb986f/?fields=*" -H "accept:
application/hal+json"
```

The response:

```

{
  "uuid": "4475a2c8-f8a0-11e8-8d33-005056bb986f",
  "name": "Cluster",
  "ipspace": {
    "uuid": "3e518ed5-f8a0-11e8-8d33-005056bb986f",
    "name": "Cluster",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/3e518ed5-f8a0-11e8-8d33-005056bb986f"
      }
    }
  }
},
"ports": [

```



```

{
  "uuid": "3e539a62-f8a0-11e8-8d33-005056bb986f",
  "name": "e0a",
  "node": {
    "name": "examplecluster-node01"
  },
  "_links": {
    "self": {
      "href": "/api/network/ethernet/ports/3e539a62-f8a0-11e8-8d33-005056bb986f"
    }
  }
},
{
  "uuid": "3e53c94a-f8a0-11e8-8d33-005056bb986f",
  "name": "e0b",
  "node": {
    "name": "examplecluster-node01"
  },
  "_links": {
    "self": {
      "href": "/api/network/ethernet/ports/3e53c94a-f8a0-11e8-8d33-005056bb986f"
    }
  }
}
],
"mtu": 9000,
"_links": {
  "self": {
    "href": "/api/network/ethernet/broadcast-domains/4475a2c8-f8a0-11e8-8d33-005056bb986f/"
  }
}
}
}
----

'''

```

=== Retrieving all broadcast domains with a specific name

The following output shows the response returned when broadcast domains with a specific name in any IPspace are requested.

```
'''
```

```

-----

# The API:
/api/network/ethernet/broadcast-domains

# The call:
curl -X GET "https://10.224.87.121/api/network/ethernet/broadcast-
domains/?name=bd1" -H "accept: application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "66b607e5-4bee-11e9-af6a-005056bb13c0",
      "name": "bd1",
      "_links": {
        "self": {
          "href": "/api/network/ethernet/broadcast-domains/66b607e5-4bee-
11e9-af6a-005056bb13c0"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/network/ethernet/broadcast-domains/?name=bd1"
    }
  }
}

```

'''

=== Retrieving the broadcast domains for an IPspace

The following output shows the response returned when the broadcast domains for a specified IPspace are requested.

'''

```

# The API:
/api/network/ethernet/broadcast-domains

```

```

# The call:

```

```

curl -X GET "https://10.224.87.121/api/network/ethernet/broadcast-
domains/?ipspace.name=Cluster&fields=*" -H "accept: application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "ae69070c-4bed-11e9-af6a-005056bb13c0",
      "name": "Cluster",
      "ipspace": {
        "uuid": "ac466a88-4bed-11e9-af6a-005056bb13c0",
        "name": "Cluster",
        "_links": {
          "self": {
            "href": "/api/network/ipspaces/ac466a88-4bed-11e9-af6a-
005056bb13c0"
          }
        }
      },
      "ports": [
        {
          "uuid": "acd67884-4bed-11e9-af6a-005056bb13c0",
          "name": "e0a",
          "node": {
            "name": "examplecluster-node-1"
          },
          "_links": {
            "self": {
              "href": "/api/network/ethernet/ports/acd67884-4bed-11e9-af6a-
005056bb13c0"
            }
          }
        },
        {
          "uuid": "acela36f-4bed-11e9-af6a-005056bb13c0",
          "name": "e0b",
          "node": {
            "name": "examplecluster-node-1"
          },
          "_links": {
            "self": {
              "href": "/api/network/ethernet/ports/acela36f-4bed-11e9-af6a-
005056bb13c0"
            }
          }
        }
      ]
    }
  ]
}

```

```

    ],
    "mtu": 1500,
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/ae69070c-4bed-
11e9-af6a-005056bb13c0"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ethernet/broadcast-
domains/?ipspace.name=Cluster&fields=*"
  }
}
}
}

```

'''

== Creating network Ethernet broadcast domains

You can use the POST API to create broadcast domains.

'''

== Example

=== Creating a new broadcast domain

The following example shows how to create a broadcast domain with a name of 'bd1' and an MTU of 1500.

'''

The API:

/api/network/ethernet/broadcast-domains

The call:

```

curl -X POST "https://<mgmt-ip>/api/network/ethernet/broadcast-
domains?return_records=true" -H "accept: application/hal+json" -d '{
"name": "bd1", "mtu": 1500 }'

```

```

# The response:
{
  "num_records": 1,
  "records": [
    {
      "name": "bd1",
      "mtu": 1500,
      "_links": {
        "self": {
          "href": "/api/network/ethernet/broadcast-domains/"
        }
      }
    }
  ]
}
-----

'''

== Updating network Ethernet broadcast domains

You can use the PATCH API to update the attributes of broadcast domains.

'''

== Examples

=== Updating the name and MTU of a specific broadcast domain

The following example shows how the PATCH request changes the broadcast
domain name to 'bd2' and the broadcast domain MTU to 9000.

'''

-----

# The API:
/api/network/ethernet/broadcast-domains/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/broadcast-
domains/6cde03b2-f8a2-11e8-8d33-005056bb986f/" -d '{ "name": "bd2", "mtu":
9000 }'
{
}
-----

```

```
'''
```

```
== Updating the IPspace of a specific broadcast domain
```

The following example shows how the PATCH request changes the IPspace of a broadcast domain to 'ipspace2'.

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ethernet/broadcast-domains/{uuid}
```

```
# The call:
```

```
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/broadcast-  
domains/c6fe2541-61f4-11e9-a66e-005056bbe83e" -d '{ "ipspace" : { "name" :  
"ipspace2" } }'
```

```
{  
}
```

```
----
```

```
'''
```

```
== Deleting network Ethernet broadcast domains
```

You can use the DELETE API to delete a broadcast domain from the cluster configuration.

```
== Example
```

```
== Deleting a specific broadcast domain
```

The following DELETE request deletes a broadcast domain.

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ethernet/broadcast-domains/{uuid}
```

```
# The call:
```

```
curl -X DELETE "https://<mgmt-ip>/api/network/ethernet/broadcast-  
domains/6cde03b2-f8a2-11e8-8d33-005056bb986f/"
```

```
----
```

```

[[ID7b590df27f6b79020a0e8c388613e904]]
= Retrieve broadcast domains for the entire cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/broadcast-domains`#

*Introduced In:* 9.6

Retrieves a collection of broadcast domains for the entire cluster.

== Related ONTAP commands

* `network port broadcast-domain show`

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|ports.name
|string
|query
|False
a|Filter by ports.name

|ports.node.name
|string
|query
|False
a|Filter by ports.node.name

|ports.uuid

```

```
|string
|query
|False
a|Filter by ports.uuid
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|ipospace.name
|string
|query
|False
a|Filter by ipospace.name
```

```
|ipospace.uuid
|string
|query
|False
a|Filter by ipospace.uuid
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|mtu
|integer
|query
|False
a|Filter by mtu
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```



```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|

|records
|array[link:#broadcast_domain[broadcast_domain]]
a|

|===

```

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ipospace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    },

```

```

        "name": "exchange",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "mtu": "1500",
    "name": "bd1",
    "ports": [
        {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "elb",
            "node": {
                "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
    ],
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====

```

```

[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description
```

```
|next
|link:href[href]
a|
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#ipospace]
[.api-collapsible-fifth-title]
ipospace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|_links
```

```

|link:#_links[_links]
a|

|name
|string
a|IPspace name

|uuid
|string
a|IPspace UUID

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of node on which the port is located.

|===

[#ports]
[.api-collapsible-fifth-title]
ports

Port UUID along with readable names

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#broadcast_domain]
[.api-collapsible-fifth-title]
broadcast_domain
```

Set of ports that will receive a broadcast Ethernet packet from any of them

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|ipSPACE
|link:#ipSPACE[ipSPACE]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
is supplied on input.
```

```
|mtu
|integer
a|Maximum transmission unit, largest packet size on this network
```

```
|name
|string
a|Name of the broadcast domain, scoped to its IPspace
```

```
|ports
|array[link:#ports[ports]]
a|Ports that belong to the broadcast domain
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====


[[ID38a2b2c2e0890ded91780699a594e966]]
= Create a new broadcast domain


[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/network/ethernet/broadcast-domains`#


*Introduced In:* 9.6


Creates a new broadcast domain.


== Required properties


* `name` - Name of the broadcast-domain to create.
* `mtu` - Maximum transmission unit (MTU) of the broadcast domain.

```

== Recommended optional properties

* `ipspace.name` or `ipspace.uuid` - IPspace the broadcast domain belongs to.

== Default property values

If not specified in POST, the following default property values are assigned:

* `ipspace` - `_Default_`

== Related ONTAP commands

* `network port broadcast-domain create`

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

```

|Description

|ipospace
|link:#ipospace[ipospace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
is supplied on input.

|mtu
|integer
a|Maximum transmission unit, largest packet size on this network

|name
|string
a|Name of the broadcast domain, scoped to its IPspace

|ports
|array[link:#ports[ports]]
a|Ports that belong to the broadcast domain

|uuid
|string
a|Broadcast domain UUID

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "ipospace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "mtu": "1500",
  "name": "bd1",
  "ports": [
    {
      "name": "e1b",
      "node": {
        "name": "node1"
      }
    }
  ]
}

```

```

    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

```

== Response

Status: 201, Created

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 1377267
| The specified IPspace does not exist.
|
| 1967082
| The specified ipspace.name does not match the IPspace name of
| ipspace.uuid.
|
| 1967102
| A POST operation might have left the configuration in an inconsistent
| state. Check the configuration.
|
| 53281982
| The specified broadcast domain name is reserved by the system.
|====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]

```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#ipospace]
```

```
[.api-collapsible-fifth-title]
```

```
ipospace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|IPspace name
```

```
|uuid
```

```
|string
```

```
a|IPspace UUID
```

```
|===
```

```
[#node]
```

```

[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of node on which the port is located.

|===

[#ports]
[.api-collapsible-fifth-title]
ports

Port UUID along with readable names

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|node
|link:#node[node]
a|

|uuid
|string
a|

|===

[#broadcast_domain]
[.api-collapsible-fifth-title]

```

broadcast_domain

Set of ports that will receive a broadcast Ethernet packet from any of them

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|ipspace
```

```
|link:#ipspace[ipspace]
```

a|Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
|mtu
```

```
|integer
```

a|Maximum transmission unit, largest packet size on this network

```
|name
```

```
|string
```

a|Name of the broadcast domain, scoped to its IPspace

```
|ports
```

```
|array[link:#ports[ports]]
```

a|Ports that belong to the broadcast domain

```
|uuid
```

```
|string
```

a|Broadcast domain UUID

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

error_arguments

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```



```

|===

//end collapsible .Definitions block
=====

[[ID356e24dcb183976e4911e4ad65e47c7b]]
= Delete a broadcast domain

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/network/ethernet/broadcast-domains/{uuid}`#

*Introduced In:* 9.6

Deletes a broadcast domain.

== Related ONTAP commands

* `network port broadcast-domain delete`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Broadcast domain UUID

|===

== Response

```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===  
| Error Code | Description  
  
| 1967103  
| A broadcast domain with ports cannot be deleted.  
|===
```

```
[cols=3*,options=header]
```

```
|===  
|Name  
|Type  
|Description
```

```
|error  
|link:#error[error]  
a|
```

```
|===
```

.Example error

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{  
  "error": {  
    "arguments": [  
      {  
        "code": "string",  
        "message": "string"  
      }  
    ],  
    "code": "4",  
    "message": "entry doesn't exist",  
    "target": "uuid"  
  }  
}
```

====

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDda845e15c8bd9f348925262243060fe3]]
= Retrieve broadcast domain details

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/broadcast-domains/{uuid}`#

*Introduced In:* 9.6

Retrieves details of a broadcast domain.

== Related ONTAP commands

* `network port broadcast-domain show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required

```

```

|Description

|uuid
|string
|path
|True
a|Broadcast domain UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
is supplied on input.

|mtu
|integer
a|Maximum transmission unit, largest packet size on this network

|name
|string
a|Name of the broadcast domain, scoped to its IPspace

```

```
|ports
|array[link:#ports[ports]]
a|Ports that belong to the broadcast domain
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "mtu": "1500",
  "name": "bd1",
  "ports": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "e1b",
      "node": {
        "name": "node1"
      }
    },

```

```

        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
],
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]

```

```
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#ipospace]
[.api-collapsible-fifth-title]
ipospace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name

|uuid
|string
a|IPspace UUID

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of node on which the port is located.

|===

[#ports]
[.api-collapsible-fifth-title]
ports

Port UUID along with readable names

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|node
|link:#node[node]
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDba106886ad467f037a7321059cd74021]]
= Update broadcast domain properties

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/network/ethernet/broadcast-domains/{uuid}`#

*Introduced In:* 9.6

```

Updates the properties of a broadcast domain.

== Related ONTAP commands

- * `network port broadcast-domain modify`
- * `network port broadcast-domain rename`
- * `network port broadcast-domain move`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Broadcast domain UUID

|===

== Request Body

[cols=3*,options=header]
|===

|Name
|Type
|Description

|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

|mtu
|integer
a|Maximum transmission unit, largest packet size on this network

```
|name
|string
a|Name of the broadcast domain, scoped to its IPspace
```

```
|ports
|array[link:#ports[ports]]
a|Ports that belong to the broadcast domain
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

.Example request

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "ipspace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "mtu": "1500",
  "name": "bd1",
  "ports": [
    {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
=====
```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

1377267	
---------	--

	The specified IPspace does not exist.
--	---------------------------------------

1377269	
---------	--

	Failed to lookup the specified IPspace.
--	-----------------------------------------

1377560	
---------	--

	Broadcast domain already exists in specified IPspace.
--	-------------------------------------------------------

1377605	
---------	--

	Moving the system-generated broadcast domain to another IPspace is not supported.
--	-----------------------------------------------------------------------------------

1967082	
---------	--

	The specified ipspace.name does not match the IPspace name of ipspace.uuid.
--	-----------------------------------------------------------------------------

1967150	
---------	--

	The specified ipspace.uuid is not valid.
--	------------------------------------------

1967151	
---------	--

	The specified ipspace.uuid and ipspace.name do not match.
--	-----------------------------------------------------------

1967152	
---------	--

	Patching IPspace for a broadcast domain requires an effective cluster version of 9.7 or later.
--	------------------------------------------------------------------------------------------------

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#ipspace]
```

```
[.api-collapsible-fifth-title]
```

```
ipspace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|IPspace name
```

```
|uuid
```

```
|string
```

```
a|IPspace UUID
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of node on which the port is located.
```

```
|===
```

```
[#ports]
[.api-collapsible-fifth-title]
ports

Port UUID along with readable names
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|node
|link:#node[node]
a|

|uuid
|string
a|

|===
```



```
[#broadcast_domain]
[.api-collapsible-fifth-title]
broadcast_domain
```

Set of ports that will receive a broadcast Ethernet packet from any of them

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|ipSPACE
|link:#ipSPACE[ipSPACE]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
is supplied on input.
```

```
|mtu
|integer
a|Maximum transmission unit, largest packet size on this network
```

```
|name
|string
a|Name of the broadcast domain, scoped to its IPspace
```

```
|ports
|array[link:#ports[ports]]
a|Ports that belong to the broadcast domain
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

:leveloffset: -1

= Manage network Ethernet ports

:leveloffset: +1

[[ID583be69994ab89d34be2cbb9831d1176]]

= Network Ethernet ports endpoint overview

== Overview

A port is a physical or virtual Ethernet network device. Physical ports may be combined into Link Aggregation Groups (LAGs or ifgrps), or divided into Virtual LANs (VLANs).

GET (collection), GET (instance), and PATCH APIs are available for all port types. POST and DELETE APIs are available for "lag" (ifgrp) and "vlan" port types.

== Retrieving network port information

The network ports GET API retrieves and displays relevant information pertaining to the ports configured in the cluster. The API retrieves the list of all ports configured in the cluster, or specifically requested ports. The fields returned in the response vary for different ports and configurations.

== Examples

=== Retrieving all ports in the cluster

The following output displays the UUID, name, and port type for all ports

configured in a 2-node cluster. The port types are physical, vlan, and lag (ifgrp).

'''

The API:

/api/network/ethernet/ports

The call:

```
curl -X GET "https://<mgmt-  
ip>/api/network/ethernet/ports?fields=uuid,name,type" -H "accept:  
application/hal+json"
```

The response:

```
{  
  "records": [  
    {  
      "uuid": "2d2c90c0-f70d-11e8-b145-005056bb5b8e",  
      "name": "e0a",  
      "type": "physical",  
      "_links": {  
        "self": {  
          "href": "/api/network/ethernet/ports/2d2c90c0-f70d-11e8-b145-  
005056bb5b8e"  
        }  
      }  
    },  
    {  
      "uuid": "2d3004da-f70d-11e8-b145-005056bb5b8e",  
      "name": "e0b",  
      "type": "physical",  
      "_links": {  
        "self": {  
          "href": "/api/network/ethernet/ports/2d3004da-f70d-11e8-b145-  
005056bb5b8e"  
        }  
      }  
    },  
    {  
      "uuid": "2d34a2cb-f70d-11e8-b145-005056bb5b8e",  
      "name": "e0c",  
      "type": "physical",  
      "_links": {  
        "self": {
```

```

        "href": "/api/network/ethernet/ports/2d34a2cb-f70d-11e8-b145-005056bb5b8e"
    }
},
{
    "uuid": "2d37189f-f70d-11e8-b145-005056bb5b8e",
    "name": "e0d",
    "type": "physical",
    "_links": {
        "self": {
            "href": "/api/network/ethernet/ports/2d37189f-f70d-11e8-b145-005056bb5b8e"
        }
    }
},
{
    "uuid": "35de5d8b-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0a",
    "type": "physical",
    "_links": {
        "self": {
            "href": "/api/network/ethernet/ports/35de5d8b-f70d-11e8-abdf-005056bb7fc8"
        }
    }
},
{
    "uuid": "35de78cc-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0b",
    "type": "physical",
    "_links": {
        "self": {
            "href": "/api/network/ethernet/ports/35de78cc-f70d-11e8-abdf-005056bb7fc8"
        }
    }
},
{
    "uuid": "35dead3c-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0c",
    "type": "physical",
    "_links": {
        "self": {
            "href": "/api/network/ethernet/ports/35dead3c-f70d-11e8-abdf-005056bb7fc8"
        }
    }
}

```

```

    }
  }
},
{
  "uuid": "35deda90-f70d-11e8-abdf-005056bb7fc8",
  "name": "e0d",
  "type": "physical",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/ports/35deda90-f70d-11e8-abdf-
005056bb7fc8"
    }
  }
},
{
  "uuid": "42e25145-f97d-11e8-ade9-005056bb7fc8",
  "name": "e0c-100",
  "type": "vlan",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/ports/42e25145-f97d-11e8-ade9-
005056bb7fc8"
    }
  }
},
{
  "uuid": "569e0abd-f97d-11e8-ade9-005056bb7fc8",
  "name": "a0a",
  "type": "lag",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/ports/569e0abd-f97d-11e8-ade9-
005056bb7fc8"
    }
  }
}
],
"num_records": 10,
"_links": {
  "self": {
    "href": "/api/network/ethernet/ports?fields=uuid,name,type"
  }
}
}
}
-----

```

```
'''
```

```
=== Retrieving a specific physical port
```

The following output displays the response when a specific physical port is requested. The system returns an error when there is no port with the requested UUID. Also, the speed field is set only if the state of the port is up.

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ethernet/ports/{uuid}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/network/ethernet/ports/2d37189f-f70d-11e8-b145-005056bb5b8e?fields=*" -H "accept: application/hal+json"
```

```
# The response:
```

```
{
  "uuid": "2d37189f-f70d-11e8-b145-005056bb5b8e",
  "name": "e0d",
  "mac_address": "00:50:56:bb:62:2d",
  "type": "physical",
  "node": {
    "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
    "name": "user-cluster-01",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/faa56898-f70c-11e8-b145-005056bb5b8e"
      }
    }
  },
  "broadcast_domain": {
    "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
    "name": "Default",
    "ipspace": {
      "name": "Default"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-b145-005056bb5b8e"
      }
    }
  }
}
```

```

    }
  },
  "enabled": true,
  "state": "up",
  "mtu": 1500,
  "speed": "1000",
  "reachability": "not_repairable",
  "reachable_broadcast_domains": [
    {
      "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
      "name": "Default",
      "ipspace": {
        "name": "Default"
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-b145-005056bb5b8e"
        }
      }
    },
    {
      "uuid": "df640ccf-72c4-11ea-b31d-005056bbfb29",
      "name": "Default-1",
      "ipspace": {
        "name": "Default"
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/broadcast-domains/df640ccf-72c4-11ea-b31d-005056bbfb29"
        }
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/network/ethernet/ports/2d37189f-f70d-11e8-b145-005056bb5b8e"
    }
  }
}
----

'''

```


=== Retrieving a specific VLAN port

The following output displays the response when a specific VLAN port is requested. The system returns an error when there is no port with the requested UUID. Also, the speed field is set only if the state of the port is up.

'''

The API:

/api/network/ethernet/ports/{uuid}

The call:

```
curl -X GET "https://<mgmt-ip>/api/network/ethernet/ports/42e25145-f97d-11e8-ade9-005056bb7fc8?fields=*" -H "accept: application/hal+json"
```

The response:

```
{
  "uuid": "42e25145-f97d-11e8-ade9-005056bb7fc8",
  "name": "e0e-100",
  "mac_address": "00:50:56:bb:52:2f",
  "type": "vlan",
  "node": {
    "uuid": "6042cf47-f70c-11e8-abdf-005056bb7fc8",
    "name": "user-cluster-02",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/6042cf47-f70c-11e8-abdf-005056bb7fc8"
      }
    }
  },
  "broadcast_domain": {
    "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
    "name": "Default",
    "ipspace": {
      "name": "Default"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-b145-005056bb5b8e"
      }
    }
  }
},
```

```

"enabled": true,
"state": "up",
"mtu": 1500,
"speed": "1000",
"reachability": "ok",
"reachable_broadcast_domains": [
  {
    "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
    "name": "Default",
    "ipspace": {
      "name": "Default"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-b145-005056bb5b8e"
      }
    }
  }
],
"vlan": {
  "tag": 100,
  "base_port": {
    "uuid": "35deff03-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0e",
    "node": {
      "name": "user-cluster-02"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/35deff03-f70d-11e8-abdf-005056bb7fc8"
      }
    }
  }
},
"_links": {
  "self": {
    "href": "/api/network/ethernet/ports/42e25145-f97d-11e8-ade9-005056bb7fc8"
  }
}
}
----

'''

```

=== Retrieving a specific LAG port

The following output displays the response when a specific LAG port is requested. The system returns an error when there is no port with the requested UUID. Also, the speed and lag.active_ports fields are set only if the state of the port is up.

'''

The API:

/api/network/ethernet/ports/{uuid}

The call:

```
curl -X GET "https://<mgmt-ip>/api/network/ethernet/ports/569e0abd-f97d-11e8-ade9-005056bb7fc8?fields=*" -H "accept: application/hal+json"
```

The response:

```
{
  "uuid": "569e0abd-f97d-11e8-ade9-005056bb7fc8",
  "name": "a0a",
  "mac_address": "02:50:56:bb:7f:c8",
  "type": "lag",
  "node": {
    "uuid": "6042cf47-f70c-11e8-abdf-005056bb7fc8",
    "name": "user-cluster-02",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/6042cf47-f70c-11e8-abdf-005056bb7fc8"
      }
    }
  },
  "broadcast_domain": {
    "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
    "name": "Default",
    "ipspace": {
      "name": "Default"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-b145-005056bb5b8e"
      }
    }
  }
}
```

```

},
"enabled": true,
"state": "up",
"mtu": 1500,
"speed": "1000",
"reachability": "repairable",
"reachable_broadcast_domains": [
  {
    "uuid": "c7934b4f-691f-11ea-87fd-005056bb1ad3",
    "name": "Default",
    "ipspace": {
      "name": "Default"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/broadcast-domains/c7934b4f-691f-11ea-87fd-005056bb1ad3"
      }
    }
  }
],
"lag": {
  "mode": "singlemode",
  "distribution_policy": "mac",
  "member_ports": [
    {
      "uuid": "35df318d-f70d-11e8-abdf-005056bb7fc8",
      "name": "e0f",
      "node": {
        "name": "user-cluster-02"
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/ports/35df318d-f70d-11e8-abdf-005056bb7fc8"
        }
      }
    },
    {
      "uuid": "35df5bad-f70d-11e8-abdf-005056bb7fc8",
      "name": "e0g",
      "node": {
        "name": "user-cluster-02"
      },
      "_links": {
        "self": {

```

```

        "href": "/api/network/ethernet/ports/35df5bad-f70d-11e8-abdf-
005056bb7fc8"
    }
}
},
{
    "uuid": "35df9926-f70d-11e8-abdf-005056bb7fc8",
    "name": "e0h",
    "node": {
        "name": "user-cluster-02"
    },
    "_links": {
        "self": {
            "href": "/api/network/ethernet/ports/35df9926-f70d-11e8-abdf-
005056bb7fc8"
        }
    }
}
],
"active_ports": [
    {
        "uuid": "35df318d-f70d-11e8-abdf-005056bb7fc8",
        "name": "e0f",
        "_links": {
            "self": {
                "href": "/api/network/ethernet/ports/35df318d-f70d-11e8-abdf-
005056bb7fc8"
            }
        }
    }
]
},
"_links": {
    "self": {
        "href": "/api/network/ethernet/ports/569e0abd-f97d-11e8-ade9-
005056bb7fc8"
    }
}
}
}
----

'''

```

=== Retrieving all LAG (ifgrp) ports in the cluster

This command retrieves all LAG ports in the cluster (that is, all ports

with type=LAG). The example shows how to filter a GET collection based on type.

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ethernet/ports
```

```
# The call:
```

```
curl -X GET "https://<mgmt-  
ip>/api/network/ethernet/ports?type=lag&node.name=user-cluster-  
01&fields=name,enabled,speed,mtu" -H "accept: application/hal+json"
```

```
# The response:
```

```
{  
  "records": [  
    {  
      "uuid": "0c226db0-4b63-11e9-8113-005056bbe040",  
      "name": "a0b",  
      "type": "lag",  
      "node": {  
        "name": "user-cluster-01"  
      },  
      "enabled": true,  
      "mtu": 1500,  
      "speed": "1000",  
      "_links": {  
        "self": {  
          "href": "/api/network/ethernet/ports/0c226db0-4b63-11e9-8113-  
005056bbe040"  
        }  
      }  
    },  
    {  
      "uuid": "d3a84153-4b3f-11e9-a00d-005056bbe040",  
      "name": "a0a",  
      "type": "lag",  
      "node": {  
        "name": "user-cluster-01"  
      },  
      "enabled": true,  
      "mtu": 1500,  
      "speed": "1000",  
      "_links": {
```

```

    "self": {
      "href": "/api/network/ethernet/ports/d3a84153-4b3f-11e9-a00d-
005056bbe040"
    }
  }
},
"num_records": 2,
"_links": {
  "self": {
    "href":
"/api/network/ethernet/ports?fields=name,enabled,speed,mtu&type=lag&node.n
ame=user-cluster-01"
  }
}
}
}
-----
'''

```

== Creating VLAN and LAG ports

You can use the network ports POST API to create VLAN and LAG ports. If you supply the optional broadcast domain property, the specified broadcast domain will be assigned to the new port immediately. Otherwise, within a few minutes automatic probing will determine the correct broadcast domain and will assign it to the port. During that period of time, the port will not be capable of hosting interfaces.

'''

== Examples

=== Creating a VLAN port

The following output displays the record returned after the creation of a VLAN port on "e0e" and VLAN tag "100".

'''

```

# The API:
/api/network/ethernet/ports

```

```

# The call:
curl -X POST "https://<mgmt-

```

```
ip>/api/network/ethernet/ports?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"type\":
\"vlan\", \"node\": { \"name\": \"user-cluster-01\" }, \"enabled\": true,
\"vlan\": { \"tag\": 100, \"base_port\": { \"name\": \"e0e\", \"node\": {
\"name\": \"user-cluster-01\" } } } } }
```

The response:

```
{
  "num_records": 1,
  "records": [
    {
      "uuid": "88b2f682-fa42-11e8-a6d7-005056bb5b8e",
      "type": "vlan",
      "node": {
        "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
        "name": "user-cluster-01",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/faa56898-f70c-11e8-b145-
005056bb5b8e"
          }
        }
      },
      "enabled": true,
      "vlan": {
        "tag": 100,
        "base_port": {
          "uuid": "2d39df72-f70d-11e8-b145-005056bb5b8e",
          "name": "e0e",
          "node": {
            "name": "user-cluster-01"
          },
          "_links": {
            "self": {
              "href": "/api/network/ethernet/ports/2d39df72-f70d-11e8-b145-
005056bb5b8e"
            }
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/ports/88b2f682-fa42-11e8-a6d7-
005056bb5b8e"
        }
      }
    }
  ]
}
```



```

    }
]
}
----

'''

=== Creating a VLAN port in a specific broadcast domain

The following output displays the record returned after the creation of a
VLAN port on "e0e" and VLAN tag "100". Also, the VLAN port is added to the
"Default" broadcast domain in the "Default" IPspace.

'''

----

# The API:
/api/network/ethernet/ports

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ethernet/ports?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
  \"type\": \"vlan\", \"node\": { \"name\": \"user-cluster-01\" },
  \"broadcast_domain\": { \"name\": \"Default\", \"ipspace\": { \"name\":
  \"Default \" } }, \"enabled\": true, \"vlan\": { \"tag\": 100,
  \"base_port\": { \"name\": \"e0e\", \"node\": { \"name\": \"user-cluster-
  01\" } } } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "88b2f682-fa42-11e8-a6d7-005056bb5b8e",
      "type": "vlan",
      "node": {
        "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
        "name": "user-cluster-01",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/faa56898-f70c-11e8-b145-
005056bb5b8e"
          }
        }
      }
    }
  ]
}

```

```

    },
    "broadcast_domain": {
      "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
      "name": "Default",
      "ipspace": {
        "name": "Default"
      },
      "_links": {
        "self": {
          "href": "/api/network/ethernet/broadcast-domains/36434bec-f70d-11e8-b145-005056bb5b8e"
        }
      }
    },
    "enabled": true,
    "vlan": {
      "tag": 100,
      "base_port": {
        "uuid": "2d39df72-f70d-11e8-b145-005056bb5b8e",
        "name": "e0e",
        "node": {
          "name": "user-cluster-01"
        },
        "_links": {
          "self": {
            "href": "/api/network/ethernet/ports/2d39df72-f70d-11e8-b145-005056bb5b8e"
          }
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/88b2f682-fa42-11e8-a6d7-005056bb5b8e"
      }
    }
  }
]
}
----

'''

```

=== Creating a LAG (ifgrp) port

The following output displays the record returned after the creation of a LAG port with "e0f", "e0g" and "e0h" as member ports.

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ethernet/ports
```

```
# The call:
```

```
curl -X POST "https://<mgmt-  
ip>/api/network/ethernet/ports?return_records=true" -H "accept:  
application/json" -H "Content-Type: application/json" -d "{ \"type\":  
\"lag\", \"node\": { \"name\": \"user-cluster-01\" }, \"enabled\": true,  
\"lag\": { \"mode\": \"singlemode\", \"distribution_policy\": \"mac\",  
\"member_ports\": [ { \"name\": \"e0f\", \"node\": { \"name\": \"user-  
cluster-01\" } }, { \"name\": \"e0g\", \"node\": { \"name\": \"user-  
cluster-01\" } }, { \"name\": \"e0h\", \"node\": { \"name\": \"user-  
cluster-01\" } } ] } }"
```

```
# The response:
```

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "uuid": "1807772a-fa4d-11e8-a6d7-005056bb5b8e",  
      "type": "lag",  
      "node": {  
        "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",  
        "name": "user-cluster-01"  
      },  
      "enabled": true,  
      "lag": {  
        "mode": "singlemode",  
        "distribution_policy": "mac",  
        "member_ports": [  
          {  
            "uuid": "2d3c9adc-f70d-11e8-b145-005056bb5b8e",  
            "name": "e0f",  
            "node": {  
              "name": "user-cluster-01"  
            }  
          },  
          {  
            "uuid": "2d40b097-f70d-11e8-b145-005056bb5b8e",
```

```

        "name": "e0g",
        "node": {
            "name": "user-cluster-01"
        }
    },
    {
        "uuid": "2d46d01e-f70d-11e8-b145-005056bb5b8e",
        "name": "e0h",
        "node": {
            "name": "user-cluster-01"
        }
    }
]
}
]
}
----
'''

```

=== Creating a LAG (ifgrp) port in a specific broadcast domain

The following output displays the record returned after the creation of a LAG port with "e0f", "e0g" and "e0h" as member ports. Also, the LAG port is added to the "Default" broadcast domain in the "Default" IPspace.

```

'''

----

# The API:
/api/network/ethernet/ports

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ethernet/ports?return_records=true" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"type\":
\"lag\", \"node\": { \"name\": \"user-cluster-01\" },
\"broadcast_domain\": { \"name\": \"Default\", \"ipspace\": { \"name\":
\"Default\" } }, \"enabled\": true, \"lag\": { \"mode\": \"singlemode\",
\"distribution_policy\": \"mac\", \"member_ports\": [ { \"name\": \"e0f\",
\"node\": { \"name\": \"user-cluster-01\" } }, { \"name\": \"e0g\",
\"node\": { \"name\": \"user-cluster-01\" } }, { \"name\": \"e0h\",
\"node\": { \"name\": \"user-cluster-01\" } } ] } }" -u admin:netapp1! -k

```

```
# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "1807772a-fa4d-11e8-a6d7-005056bb5b8e",
      "type": "lag",
      "node": {
        "uuid": "faa56898-f70c-11e8-b145-005056bb5b8e",
        "name": "user-cluster-01"
      },
      "broadcast_domain": {
        "uuid": "36434bec-f70d-11e8-b145-005056bb5b8e",
        "name": "Default",
        "ipspace": {
          "name": "Default"
        }
      },
      "enabled": true,
      "lag": {
        "mode": "singlemode",
        "distribution_policy": "mac",
        "member_ports": [
          {
            "uuid": "2d3c9adc-f70d-11e8-b145-005056bb5b8e",
            "name": "e0f",
            "node": {
              "name": "user-cluster-01"
            }
          },
          {
            "uuid": "2d40b097-f70d-11e8-b145-005056bb5b8e",
            "name": "e0g",
            "node": {
              "name": "user-cluster-01"
            }
          },
          {
            "uuid": "2d46d01e-f70d-11e8-b145-005056bb5b8e",
            "name": "e0h",
            "node": {
              "name": "user-cluster-01"
            }
          }
        ]
      }
    }
  ]
}
```

```

    }
]
}
----

'''

== Updating ports

You can use the network ports PATCH API to update the attributes of ports.

'''

== Examples

=== Updating the broadcast domain of a port

The following PATCH request removes the port from the current broadcast
domain and adds it to the specified broadcast domain.

'''

----

# The API:
/api/network/ethernet/ports/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/ports/6867efaf-d702-
11e8-994f-005056bbc994" -H "accept: application/hal+json" -H "Content-
Type: application/json" -d "{ \"broadcast_domain\": { \"name\":
\"Default\", \"ipspace\": { \"name\": \"Default\" } } }"
----

'''

=== Updating the admin status of a port

The following PATCH request brings the specified port down.

'''

----

# The API:
/api/network/ethernet/ports/{uuid}

```

```
# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/ports/51d3ab39-d86d-11e8-aca6-005056bbc994" -H "accept: application/hal+json" -H "Content-Type: application/json" -d "{ \"enabled\": \"false\" }"
----
```

'''

=== Repairing a port

The following PATCH request repairs a port. Only ports that have reachability as "repairable" can be repaired. The "reachability" parameter cannot be patched in the same request as other parameters that might affect the target port's reachability status.

'''

```
# The API:
/api/network/ethernet/ports/{uuid}
```

```
# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/ports/51d3ab39-d86d-11e8-aca6-005056bbc994" -H "accept: application/hal+json" -H "Content-Type: application/json" -d "{ \"reachability\": \"ok\" }"
----
```

'''

== Deleting ports

You can use the network ports DELETE API to delete VLAN and LAG ports in the cluster. Note that physical ports cannot be deleted. Deleting a port also removes the port from the broadcast domain.

'''

== Example

=== Deleting a VLAN port

The network ports DELETE API is used to delete a VLAN port.

'''

```

# The API:
/api/network/ethernet/ports/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ethernet/ports/6867efaf-
d702-11e8-994f-005056bbc994" -H "accept: application/hal+json" -H
"Content-Type: application/json"
----

'''

[[ID6e7cdac996b6879c510234ae0b086e86]]
= Retrieve ports

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/ports`#

*Introduced In:* 9.6

Retrieves a collection of ports (physical, VLAN and LAG) for an entire
cluster.

== Related ONTAP commands

* `network port show`
* `network port ifgrp show`
* `network port vlan show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|reachable_broadcast_domains.uuid
|string

```



```
|query
|False
a|Filter by reachable_broadcast_domains.uuid
```

* Introduced in: 9.8

```
|reachable_broadcast_domains.name
|string
|query
|False
a|Filter by reachable_broadcast_domains.name
```

* Introduced in: 9.8

```
|reachable_broadcast_domains.ipspace.name
|string
|query
|False
a|Filter by reachable_broadcast_domains.ipspace.name
```

* Introduced in: 9.8

```
|state
|string
|query
|False
a|Filter by state
```

```
|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp
```

* Introduced in: 9.8

```
|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write
```

* Introduced in: 9.8

```
|statistics.throughput_raw.read  
|integer  
|query  
|False  
a|Filter by statistics.throughput_raw.read
```

* Introduced in: 9.8

```
|statistics.throughput_raw.total  
|integer  
|query  
|False  
a|Filter by statistics.throughput_raw.total
```

* Introduced in: 9.8

```
|statistics.device.receive_raw.packets  
|integer  
|query  
|False  
a|Filter by statistics.device.receive_raw.packets
```

* Introduced in: 9.8

```
|statistics.device.receive_raw.discards  
|integer  
|query  
|False  
a|Filter by statistics.device.receive_raw.discards
```

* Introduced in: 9.8

```
|statistics.device.receive_raw.errors  
|integer  
|query  
|False  
a|Filter by statistics.device.receive_raw.errors
```

* Introduced in: 9.8

```
|statistics.device.timestamp
|string
|query
|False
a|Filter by statistics.device.timestamp
```

* Introduced in: 9.8

```
|statistics.device.transmit_raw.packets
|integer
|query
|False
a|Filter by statistics.device.transmit_raw.packets
```

* Introduced in: 9.8

```
|statistics.device.transmit_raw.discards
|integer
|query
|False
a|Filter by statistics.device.transmit_raw.discards
```

* Introduced in: 9.8

```
|statistics.device.transmit_raw.errors
|integer
|query
|False
a|Filter by statistics.device.transmit_raw.errors
```

* Introduced in: 9.8

```
|statistics.device.link_down_count_raw
|integer
|query
|False
a|Filter by statistics.device.link_down_count_raw
```

* Introduced in: 9.8

```
|statistics.status
```

```
|string
|query
|False
a|Filter by statistics.status
```

* Introduced in: 9.8

```
|speed
|integer
|query
|False
a|Filter by speed
```

```
|reachability
|string
|query
|False
a|Filter by reachability
```

* Introduced in: 9.8

```
|mac_address
|string
|query
|False
a|Filter by mac_address
```

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|node.name
|string
|query
|False
a|Filter by node.name
```

```
|node.uuid
|string
```

```
|query
|False
a|Filter by node.uuid
```

```
|mtu
|integer
|query
|False
a|Filter by mtu
```

```
|vlan.tag
|integer
|query
|False
a|Filter by vlan.tag
```

```
|vlan.base_port.name
|string
|query
|False
a|Filter by vlan.base_port.name
```

```
|vlan.base_port.node.name
|string
|query
|False
a|Filter by vlan.base_port.node.name
```

```
|vlan.base_port.uuid
|string
|query
|False
a|Filter by vlan.base_port.uuid
```

```
|broadcast_domain.uuid
|string
|query
|False
a|Filter by broadcast_domain.uuid
```

```
|broadcast_domain.name
|string
|query
|False
a|Filter by broadcast_domain.name

|broadcast_domain.ipspace.name
|string
|query
|False
a|Filter by broadcast_domain.ipspace.name

|lag.active_ports.name
|string
|query
|False
a|Filter by lag.active_ports.name

|lag.active_ports.node.name
|string
|query
|False
a|Filter by lag.active_ports.node.name

|lag.active_ports.uuid
|string
|query
|False
a|Filter by lag.active_ports.uuid

|lag.mode
|string
|query
|False
a|Filter by lag.mode

|lag.member_ports.name
|string
|query
|False
a|Filter by lag.member_ports.name
```

```
|lag.member_ports.node.name
|string
|query
|False
a|Filter by lag.member_ports.node.name
```

```
|lag.member_ports.uuid
|string
|query
|False
a|Filter by lag.member_ports.uuid
```

```
|lag.distribution_policy
|string
|query
|False
a|Filter by lag.distribution_policy
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|metric.timestamp
|string
|query
|False
a|Filter by metric.timestamp
```

* Introduced in: 9.8

```
|metric.duration
```

```
|string
|query
|False
a|Filter by metric.duration
```

* Introduced in: 9.8

```
|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write
```

* Introduced in: 9.8

```
|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read
```

* Introduced in: 9.8

```
|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total
```

* Introduced in: 9.8

```
|metric.status
|string
|query
|False
a|Filter by metric.status
```

* Introduced in: 9.8

```
|type
|string
|query
```



```

|False
a|Filter by type

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.

```

Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|

|records

|array[link:#port[port]]

a|

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "_links": {

 "next": {

 "href": "/api/resourcelink"

 },

 "self": {

 "href": "/api/resourcelink"

 }

 },

 "records": [

 {

 "_links": {

 "self": {

```

        "href": "/api/resourcelink"
    }
},
"broadcast_domain": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "ipspace": {
        "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"lag": {
    "active_ports": [
        {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "elb",
            "node": {
                "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
    ],
    "distribution_policy": "string",
    "member_ports": [
        {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "elb",
            "node": {
                "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
    ],
    "mode": "string"

```

```

},
"mac_address": "01:02:03:04:05:06",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"mtu": "1500",
"name": "elb",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"reachability": "ok",
"reachable_broadcast_domains": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"speed": "1000",
"state": "string",
"statistics": {

```

```

    "device": {
      "link_down_count_raw": "3",
      "receive_raw": {
        "discards": "100",
        "errors": "200",
        "packets": "500"
      },
      "timestamp": "2017-01-25T11:20:13Z",
      "transmit_raw": {
        "discards": "100",
        "errors": "200",
        "packets": "500"
      }
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "type": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vlan": {
    "base_port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "elb",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "tag": "100"
  }
}

====

== Error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

```

```

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the broadcast domain's IPspace

```

```

|===

[#broadcast_domain]
[.api-collapsible-fifth-title]
broadcast_domain

Broadcast domain UUID along with a readable name. Either the UUID or both
names may be provided on input.

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ipspace
|link:#ipspace[ipspace]
a|

|name
|string
a|Name of the broadcast domain, scoped to its IPspace

```



```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|Name of node on which the port is located.
```

```
|===
```

```
[#active_ports]
[.api-collapsible-fifth-title]
active_ports
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
```

```
a|
```

```
|node  
|link:#node[node]
```

```
a|
```

```
|uuid  
|string
```

```
a|
```

```
|===
```

```
[#member_ports]  
[.api-collapsible-fifth-title]  
member_ports
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]
```

```
a|
```

```
|name  
|string
```

```
a|
```

```
|node  
|link:#node[node]
```

```
a|
```

```
|uuid  
|string
```

```
a|
```

```
|===
```

```
[#lag]
```

```
[.api-collapsible-fifth-title]
lag

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active_ports
|array[link:#active_ports[active_ports]]
a|Active ports of a LAG (ifgrp). (Some member ports may be inactive.)

|distribution_policy
|string
a|Policy for mapping flows to ports for outbound packets through a LAG
(ifgrp).

|member_ports
|array[link:#member_ports[member_ports]]
a|

|mode
|string
a|Determines how the ports interact with the switch.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the port object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

```
metric
```

The most recent sample of I/O metrics for the port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|duration
```

```
|string
```

```
a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
```

```
|status
```

```
|string
```

```
a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data".
```

"inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the port object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#reachable_broadcast_domains]
```

```
[.api-collapsible-fifth-title]
```

```
reachable_broadcast_domains
```

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|ipspace
```

```
|link:#ipspace[ipspace]
```

```
a|
```

```
|name
```

```
|string
```

```
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid
```

```
|string
```

```
a|Broadcast domain UUID
```

```
|===
```

```
[#receive_raw]
```

```
[.api-collapsible-fifth-title]
```

```
receive_raw
```

Packet receive counters for the Ethernet port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|discards
|integer
a|Total number of discarded packets.
```

```
|errors
|integer
a|Number of packet errors.
```

```
|packets
|integer
a|Total packet count.
```

```
|===
```

```
[#transmit_raw]
[.api-collapsible-fifth-title]
transmit_raw
```

Packet transmit counters for the Ethernet port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|discards
|integer
a|Total number of discarded packets.
```

```
|errors
|integer
a|Number of packet errors.
```

```
|packets
|integer
a|Total packet count.
```

```
|===
```

```
[#device]
[.api-collapsible-fifth-title]
device
```

Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|link_down_count_raw
|integer
```

a|The number of link state changes from up to down seen on the device.

```
|receive_raw
|link:#receive_raw[receive_raw]
```

a|Packet receive counters for the Ethernet port.

```
|timestamp
|string
```

a|The timestamp when the device specific counters were collected.

```
|transmit_raw
|link:#transmit_raw[transmit_raw]
```

a|Packet transmit counters for the Ethernet port.

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the port object. This can be used along with

delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

The real time I/O statistics for the port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|device
```

```
|link:#device[device]
```

```
a|Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two
```

samples.

|status

|string

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data".

"inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value.

"inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the throughput_raw performance data.

|===

[#base_port]

[.api-collapsible-fifth-title]

base_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

[cols=3*,options=header]

|===

|Name

|Type

|Description

```

|_links
|link:#_links[_links]
a|

|name
|string
a|

|node
|link:#node[node]
a|

|uuid
|string
a|

|===

[#vlan]
[.api-collapsible-fifth-title]
vlan

[cols=3*,options=header]
|===
|Name
|Type
|Description

|base_port
|link:#base_port[base_port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.

|tag
|integer
a|VLAN ID

|===

[#port]
[.api-collapsible-fifth-title]
port

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name. Either the UUID or
both names may be provided on input.

|enabled
|boolean
a|

|lag
|link:#lag[lag]
a|

|mac_address
|string
a|

|metric
|link:#metric[metric]
a|The most recent sample of I/O metrics for the port.

|mtu
|integer
a|MTU of the port in bytes. Set by broadcast domain.

|name
|string
a|Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-
200 (vlan on LAG/ifgrp)

|node
|link:#node[node]
a|

```

```

|reachability
|string
a|Reachability status of the port. Enum value "ok" is the only acceptable
value for a PATCH request to repair a port.

|reachable_broadcast_domains
|array[link:#reachable_broadcast_domains[reachable_broadcast_domains]]
a|Reachable broadcast domains.

|speed
|integer
a|Link speed in Mbps

|state
|string
a|Operational state of the port.

|statistics
|link:#statistics[statistics]
a|The real time I/O statistics for the port.

|type
|string
a|Type of physical or virtual port

|uuid
|string
a|Port UUID

|vlan
|link:#vlan[vlan]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[IDf9f126f8d261a7e605e434d7c8239e31]]

= Create a new VLAN or LAG

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/network/ethernet/ports`#

Introduced In: 9.6

Creates a new VLAN (such as node1:e0a-100) or LAG (ifgrp, such as node2:a0a).

== Required properties

* `node` - Node the port will be created on.

* `type` - Defines if a VLAN or LAG will be created:

*** VLAN

***** `vlan.base_port` - Physical port or LAG the VLAN will be created on.

***** `vlan.tag` - Tag used to identify VLAN on the base port.

*** LAG

***** `lag.mode` - Policy for the LAG that will be created.

***** `lag.distribution_policy` - Indicates how the packets are distributed between ports.

***** `lag.member_ports` - Set of ports the LAG consists of.

== Optional properties

* `broadcast_domain` - The layer-2 broadcast domain the port is associated with. The port will be placed in a broadcast domain if it is not specified. It may take several minutes for the broadcast domain to be

assigned. During that period the port cannot host interfaces.

== Related ONTAP commands

- * `network port ifgrp create`
- * `network port vlan create`

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

return_records
boolean
query
False

a|The default is false. If set to true, the records are returned.

- * Default value:

|===

== Request Body

[cols=3*,options=header]
|===

Name
Type
Description

broadcast_domain
link:#broadcast_domain[broadcast_domain]

a|Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

enabled
boolean


```

a|

|lag
|link:#lag[lag]
a|

|mac_address
|string
a|

|mtu
|integer
a|MTU of the port in bytes. Set by broadcast domain.

|name
|string
a|Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-
200 (vlan on LAG/ifgrp)

|node
|link:#node[node]
a|

|reachable_broadcast_domains
|array[link:#reachable_broadcast_domains[reachable_broadcast_domains]]
a|Reachable broadcast domains.

|speed
|integer
a|Link speed in Mbps

|state
|string
a|Operational state of the port.

|type
|string
a|Type of physical or virtual port

|uuid
|string

```

a|Port UUID

|vlan
|link:#vlan[vlan]
a|

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "broadcast_domain": {
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "lag": {
    "active_ports": [
      {
        "name": "elb",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "distribution_policy": "string",
    "member_ports": [
      {
        "name": "elb",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "mode": "string"
  },
  "mac_address": "01:02:03:04:05:06",
  "mtu": "1500",
```

```

"name": "e1b",
"node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"reachable_broadcast_domains": [
  {
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"speed": "1000",
"state": "string",
"type": "string",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"vlan": {
  "base_port": {
    "name": "e1b",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "tag": "100"
}
}
====

== Response

```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|

```

```
|records
|array[link:#port[port]]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "records": [
    {
      "broadcast_domain": {
        "ipspace": {
          "name": "ipspace1"
        },
        "name": "bd1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "lag": {
        "active_ports": [
          {
            "name": "e1b",
            "node": {
              "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "distribution_policy": "string",
        "member_ports": [
          {
            "name": "e1b",
            "node": {
              "name": "node1"
            },
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "mode": "string"
      },
      "mac_address": "01:02:03:04:05:06",
      "mtu": "1500",
      "name": "e1b",
```

```

"node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"reachable_broadcast_domains": [
  {
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"speed": "1000",
"state": "string",
"type": "string",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"vlan": {
  "base_port": {
    "name": "e1b",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "tag": "100"
}
]
}
=====

```

== Error

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 1376361
| Port is already a member of a LAG.

| 1966189

```

| Port is the home port or current port of an interface.

| 1967083
| The specified type is not valid.

| 1967084
| The specified node UUID is not valid.

| 1967085
| The specified node name is not valid.

| 1967086
| Node name and UUID must match if both are provided.

| 1967087
| The specified broadcast domain UUID is not valid.

| 1967088
| The specified broadcast domain name does not exist in the specified IPspace.

| 1967089
| The specified broadcast domain UUID, name, and IPspace name do not match.

| 1967090
| The specified VLAN base port UUID is not valid.

| 1967091
| The specified VLAN base port name and node name are not valid.

| 1967092
| The specified node does not match the node specified for the VLAN base port.

| 1967093
| The specified VLAN base port UUID, name, and VLAN base port node name do not match.

| 1967094
| The specified LAG member port UUID is not valid.

| 1967095
| The specified LAG member port name and node name combination is not valid.

| 1967096

```
| The specified node does not match the specified LAG member port node.

| 1967097
| The specified LAG member ports UUID, name, and node name do not match.

| 1967098
| VLAN POST operation has failed because admin status could not be set for
the specified port.

| 1967099
| Partial success of the VLAN POST operation. Verify the state of the
created VLAN for more information.

| 1967100
| LAG POST operation failed because admin status could not be set.

| 1967101
| Partial success of the LAG POST operation. Verify the state of the
created LAG for more information.

| 1967102
| POST operation might have left the configuration in an inconsistent
state. Check the configuration.

| 1967148
| Failure to remove port from broadcast domain.

| 1967149
| Failure to add port to broadcast domain.

| 1967175
| VLANs cannot be created on ports in the Cluster IPspace.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#ipspace]
[.api-collapsible-fifth-title]
ipspace

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the broadcast domain's IPspace

|===

[#broadcast_domain]
[.api-collapsible-fifth-title]
broadcast_domain

Broadcast domain UUID along with a readable name. Either the UUID or both
names may be provided on input.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```
|ipospace
|link:#ipospace[ipospace]
a|
```

```
|name
|string
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|Name of node on which the port is located.
```

```
|===
```

```
[#active_ports]
[.api-collapsible-fifth-title]
active_ports
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#member_ports]
```

```
[.api-collapsible-fifth-title]
```

```
member_ports
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#lag]
```

```
[.api-collapsible-fifth-title]
lag

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active_ports
|array[link:#active_ports[active_ports]]
a|Active ports of a LAG (ifgrp). (Some member ports may be inactive.)

|distribution_policy
|string
a|Policy for mapping flows to ports for outbound packets through a LAG
(ifgrp).

|member_ports
|array[link:#member_ports[member_ports]]
a|

|mode
|string
a|Determines how the ports interact with the switch.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the port object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

```
metric
```

The most recent sample of I/O metrics for the port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|duration
```

```
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|status
```

```
|string
```

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value.

"inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the port object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#node]

[.api-collapsible-fifth-title]

node

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|

|uuid

|string

a|

|===

[#reachable_broadcast_domains]

[.api-collapsible-fifth-title]

reachable_broadcast_domains

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

|ipospace
|link:#ipospace[ipospace]
a|

|name
|string
a|Name of the broadcast domain, scoped to its IPspace

|uuid
|string
a|Broadcast domain UUID

|===

[#receive_raw]
[.api-collapsible-fifth-title]
receive_raw

Packet receive counters for the Ethernet port.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|discards
|integer
a|Total number of discarded packets.

|errors
|integer
a|Number of packet errors.

|packets
|integer

```

```
a|Total packet count.
```

```
|===
```

```
[#transmit_raw]  
[.api-collapsible-fifth-title]  
transmit_raw
```

Packet transmit counters for the Ethernet port.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|discards  
|integer  
a|Total number of discarded packets.
```

```
|errors  
|integer  
a|Number of packet errors.
```

```
|packets  
|integer  
a|Total packet count.
```

```
|===
```

```
[#device]  
[.api-collapsible-fifth-title]  
device
```

Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|link_down_count_raw
|integer
a|The number of link state changes from up to down seen on the device.
```

```

|timestamp
|string
a|The timestamp when the device specific counters were collected.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.
```

```

|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```

|write
|integer
a|Performance metric for write I/O operations.
```



```
|===
```

```
[#statistics]  
[.api-collapsible-fifth-title]  
statistics
```

The real time I/O statistics for the port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|status  
|string
```

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw  
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp  
|string
```

a|The timestamp of the throughput_raw performance data.

```
|===
```

```
[#base_port]
[.api-collapsible-fifth-title]
base_port
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#vlan]
[.api-collapsible-fifth-title]
vlan
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|base_port
|link:#base_port[base_port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.
```

```
|tag
```

```

|integer
a|VLAN ID

|===

[#port]
[.api-collapsible-fifth-title]
port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name. Either the UUID or
both names may be provided on input.

|enabled
|boolean
a|

|lag
|link:#lag[lag]
a|

|mac_address
|string
a|

|mtu
|integer
a|MTU of the port in bytes. Set by broadcast domain.

|name
|string
a|Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-
200 (vlan on LAG/ifgrp)

|node

```

```

|link:#node[node]
a|

|reachable_broadcast_domains
|array[link:#reachable_broadcast_domains[reachable_broadcast_domains]]
a|Reachable broadcast domains.

|speed
|integer
a|Link speed in Mbps

|state
|string
a|Operational state of the port.

|type
|string
a|Type of physical or virtual port

|uuid
|string
a|Port UUID

|vlan
|link:#vlan[vlan]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string

```

```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[IDf254538b0cb3c99b9d32a4236a6c48e6]]
```

= Delete a VLAN or LAG

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/network/ethernet/ports/{uuid}`#
```

Introduced In: 9.6

Deletes a VLAN or LAG.

== Related ONTAP commands

* `network port ifgrp delete`

* `network port vlan delete`

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Port UUID
```

```
|===
```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 1376858

| Port already has an interface bound.

| 1966189

| Port is the home port or current port of an interface.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```



```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[IDaafe6d783084e7e546c08b20c0cee8e1]]
= Retrieve a physical port, VLAN, or LAG details
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/ports/{uuid}`#
```

```
*Introduced In:* 9.6
```

Retrieves the details of a physical port, VLAN, or LAG.

```
== Related ONTAP commands
```

```
* `network port show`
* `network port ifgrp show`
* `network port vlan show`
```

```
== Parameters
```

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|uuid
|string
|path
|True
a|Port UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name. Either the UUID or
both names may be provided on input.

|enabled
|boolean
a|

|lag
|link:#lag[lag]
a|

|mac_address
|string
```

```

a|

|metric
|link:#metric[metric]
a|The most recent sample of I/O metrics for the port.

|mtu
|integer
a|MTU of the port in bytes. Set by broadcast domain.

|name
|string
a|Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-
200 (vlan on LAG/ifgrp)

|node
|link:#node[node]
a|

|reachability
|string
a|Reachability status of the port. Enum value "ok" is the only acceptable
value for a PATCH request to repair a port.

|reachable_broadcast_domains
|array[link:#reachable_broadcast_domains[reachable_broadcast_domains]]
a|Reachable broadcast domains.

|speed
|integer
a|Link speed in Mbps

|state
|string
a|Operational state of the port.

|statistics
|link:#statistics[statistics]
a|The real time I/O statistics for the port.

```

```
|type
|string
a|Type of physical or virtual port
```

```
|uuid
|string
a|Port UUID
```

```
|vlan
|link:#vlan[vlan]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
=====
```

```
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "broadcast_domain": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "lag": {
    "active_ports": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    ]
  }
}
```

```

    }
    },
    "name": "elb",
    "node": {
        "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
}
],
"distribution_policy": "string",
"member_ports": [
    {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "elb",
        "node": {
            "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
    }
],
"mode": "string"
},
"mac_address": "01:02:03:04:05:06",
"metric": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "duration": "PT15S",
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"mtu": "1500",
"name": "elb",
"node": {
    "_links": {

```

```

    "self": {
      "href": "/api/resourcelink"
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "reachability": "ok",
  "reachable_broadcast_domains": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ipspace": {
        "name": "ipspace1"
      },
      "name": "bd1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "speed": "1000",
  "state": "string",
  "statistics": {
    "device": {
      "link_down_count_raw": "3",
      "receive_raw": {
        "discards": "100",
        "errors": "200",
        "packets": "500"
      },
      "timestamp": "2017-01-25T11:20:13Z",
      "transmit_raw": {
        "discards": "100",
        "errors": "200",
        "packets": "500"
      }
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  }

```

```

},
"type": "string",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"vlan": {
  "base_port": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "elb",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "tag": "100"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [

```

```

    {
      "code": "string",
      "message": "string"
    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

|===

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```

|self
|link:#href[href]
a|

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the broadcast domain's IPspace

|===

[#broadcast_domain]
[.api-collapsible-fifth-title]
broadcast_domain

Broadcast domain UUID along with a readable name. Either the UUID or both
names may be provided on input.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ipspace
|link:#ipspace[ipspace]
a|

```

```
|name
|string
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|Name of node on which the port is located.
```

```
|===
```

```
[#active_ports]
[.api-collapsible-fifth-title]
active_ports
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
```

```
a|
```

```
|name  
|string  
a|
```

```
|node  
|link:#node[node]  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#member_ports]  
[.api-collapsible-fifth-title]  
member_ports
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
 |_links  
 |link:#_links[_links]  
 a|
```

```
|name  
|string  
a|
```

```
|node  
|link:#node[node]  
a|
```

```
|uuid  
|string  
a|
```

```

|===

[#lag]
[.api-collapsible-fifth-title]
lag

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active_ports
|array[link:#active_ports[active_ports]]
a|Active ports of a LAG (ifgrp). (Some member ports may be inactive.)

|distribution_policy
|string
a|Policy for mapping flows to ports for outbound packets through a LAG
(ifgrp).

|member_ports
|array[link:#member_ports[member_ports]]
a|

|mode
|string
a|Determines how the ports interact with the switch.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the port object.

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

The most recent sample of I/O metrics for the port.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|status
|string
a|Errors associated with the sample. For example, if the aggregation of

```

data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the port object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#reachable_broadcast_domains]  
[.api-collapsible-fifth-title]  
reachable_broadcast_domains
```

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|ipspace  
|link:#ipspace[ipspace]  
a|
```

```
|name  
|string  
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid  
|string  
a|Broadcast domain UUID
```

```
|===
```

```
[#receive_raw]  
[.api-collapsible-fifth-title]  
receive_raw
```

Packet receive counters for the Ethernet port.

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|discards
|integer
a|Total number of discarded packets.

|errors
|integer
a|Number of packet errors.

|packets
|integer
a|Total packet count.

|===

[#transmit_raw]
[.api-collapsible-fifth-title]
transmit_raw

Packet transmit counters for the Ethernet port.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|discards
|integer
a|Total number of discarded packets.

|errors
|integer
a|Number of packet errors.

|packets
|integer

```



```
a|Total packet count.
```

```
|===
```

```
[#device]
```

```
[.api-collapsible-fifth-title]
```

```
device
```

Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|link_down_count_raw
```

```
|integer
```

```
a|The number of link state changes from up to down seen on the device.
```

```
|receive_raw
```

```
|link:#receive_raw[receive_raw]
```

```
a|Packet receive counters for the Ethernet port.
```

```
|timestamp
```

```
|string
```

```
a|The timestamp when the device specific counters were collected.
```

```
|transmit_raw
```

```
|link:#transmit_raw[transmit_raw]
```

```
a|Packet transmit counters for the Ethernet port.
```

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

The real time I/O statistics for the port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|device
```

```
|link:#device[device]
```

```
a|Device-related counters for the port object. These counters are
```

applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

|status

|string

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the throughput_raw performance data.

|===

[#base_port]

[.api-collapsible-fifth-title]

base_port

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

[cols=3*,options=header]

|===

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|node
|link:#node[node]
a|

|uuid
|string
a|

|===

[#vlan]
[.api-collapsible-fifth-title]
vlan

[cols=3*,options=header]
|===
|Name
|Type
|Description

|base_port
|link:#base_port[base_port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.

|tag
|integer
a|VLAN ID

|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID301123fe5de92cb7015d2943faf96b4]]
= Update a port

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/network/ethernet/ports/{uuid}`#

*Introduced In:* 9.6

Updates a port.

== Related ONTAP commands

* `network port broadcast-domain add-ports`
* `network port broadcast-domain remove-ports`
* `network port ifgrp modify`
* `network port modify`
* `network port vlan modify`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path

```

```

|True
a|Port UUID

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name. Either the UUID or
both names may be provided on input.

|enabled
|boolean
a|

|lag
|link:#lag[lag]
a|

|mac_address
|string
a|

|mtu
|integer
a|MTU of the port in bytes. Set by broadcast domain.

|name
|string
a|Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-
200 (vlan on LAG/ifgrp)

|reachability
|string
a|Reachability status of the port. Enum value "ok" is the only acceptable

```

value for a PATCH request to repair a port.

```
|reachable_broadcast_domains
|array[link:#reachable_broadcast_domains[reachable_broadcast_domains]]
a|Reachable broadcast domains.
```

```
|speed
|integer
a|Link speed in Mbps
```

```
|state
|string
a|Operational state of the port.
```

```
|uuid
|string
a|Port UUID
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "broadcast_domain": {
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "lag": {
    "active_ports": [
      {
        "name": "elb",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ]
  }
}
```



```

],
"member_ports": [
  {
    "name": "elb",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"mac_address": "01:02:03:04:05:06",
"mtu": "1500",
"name": "elb",
"reachability": "ok",
"reachable_broadcast_domains": [
  {
    "ipspace": {
      "name": "ipspace1"
    },
    "name": "bd1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"speed": "1000",
"state": "string",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response

```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 1376361
| Port is already a member of a LAG.

```

| 1377562
| Port cannot be used because it is currently the home port or current port of an interface.

| 1377563
| Port is already a member of a LAG.

| 1967087
| The specified broadcast domain UUID is not valid.

| 1967088
| The specified broadcast domain name does not exist in the specified IPspace.

| 1967089
| The specified broadcast domain UUID, name and IPspace name do not match.

| 1967094
| The specified LAG member port UUID is not valid.

| 1967095
| The specified LAG member port name and node name combination is not valid.

| 1967096
| The specified node does not match the specified LAG member port node.

| 1967097
| The specified LAG member ports UUID, name, and node name do not match.

| 1967148
| Failure to remove port from broadcast domain.

| 1967149
| Failure to add port to broadcast domain.

| 1967184
| The reachability parameter cannot be patched in the same request as other parameters that might affect the target port's reachability status.

| 1967185
| The port cannot be repaired because the port is deemed as non-repairable.

| 1967186
| Invalid value for the reachability parameter.

```
| 1967580
| This command is not supported as the effective cluster version is
earlier than 9.8.
```

```
| 1967582
| The reachability parameter is not supported on this cluster.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

====

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|href
|string
a|
```

|===

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#ipospace]
[.api-collapsible-fifth-title]
ipospace
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|name
|string
a|Name of the broadcast domain's IPspace
```

```
|===
```

```
[#broadcast_domain]
[.api-collapsible-fifth-title]
broadcast_domain
```

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|ipspace
|link:#ipspace[ipspace]
a|
```

```
|name
|string
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
```

|Description

|name

|string

a|Name of node on which the port is located.

|===

[#active_ports]

[.api-collapsible-fifth-title]

active_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|

|node

|link:#node[node]

a|

|uuid

|string

a|

|===

[#member_ports]

[.api-collapsible-fifth-title]

member_ports

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|node
|link:#node[node]
a|

|uuid
|string
a|

|===

[#lag]
[.api-collapsible-fifth-title]
lag

[cols=3*,options=header]
|===
|Name
|Type
|Description

|active_ports
|array[link:#active_ports[active_ports]]
a|Active ports of a LAG (ifgrp). (Some member ports may be inactive.)

|member_ports
|array[link:#member_ports[member_ports]]
a|

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

```

The rate of throughput bytes per second observed at the port object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

```
metric
```

The most recent sample of I/O metrics for the port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|duration
```

```
|string
```

```
a|The duration over which this sample is calculated. The time durations  
are represented in the ISO-8601 standard format. Samples can be calculated  
over the following durations:
```

```

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"inconsistent_old_data" is returned when one or more nodes do not have the
latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the port object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

```



```
[#reachable_broadcast_domains]
[.api-collapsible-fifth-title]
reachable_broadcast_domains
```

Broadcast domain UUID along with a readable name. Either the UUID or both names may be provided on input.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|ipspace
|link:#ipspace[ipspace]
a|
```

```
|name
|string
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#receive_raw]
[.api-collapsible-fifth-title]
receive_raw
```

Packet receive counters for the Ethernet port.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|discards
```

```
|integer
a|Total number of discarded packets.
```

```
|errors
|integer
a|Number of packet errors.
```

```
|packets
|integer
a|Total packet count.
```

```
|===
```

```
[#transmit_raw]
[.api-collapsible-fifth-title]
transmit_raw
```

Packet transmit counters for the Ethernet port.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|discards
|integer
a|Total number of discarded packets.
```

```
|errors
|integer
a|Number of packet errors.
```

```
|packets
|integer
a|Total packet count.
```

```
|===
```

```
[#device]
[.api-collapsible-fifth-title]
device
```

Device-related counters for the port object. These counters are applicable at the lowest layer of the networking stack. These values can be used to calculate both transmit and receive packet and error rates by comparing two samples taken at different times and calculating the increase in counter value divided by the elapsed time between the two samples.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|link_down_count_raw
|integer
a|The number of link state changes from up to down seen on the device.
```

```
|timestamp
|string
a|The timestamp when the device specific counters were collected.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the port object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

The real time I/O statistics for the port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|status
```

```
|string
```

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data".

"inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value.

"inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
```

```
|link:#throughput_raw[throughput_raw]
```

```
a|Throughput bytes observed at the port object. This can be used along
```

with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp
|string
a|The timestamp of the throughput_raw performance data.
```

```
|===
```

```
[#base_port]
[.api-collapsible-fifth-title]
base_port
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#vlan]
[.api-collapsible-fifth-title]
vlan
```

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|base_port
|link:#base_port[base_port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.

|tag
|integer
a|VLAN ID

|===

[#port]
[.api-collapsible-fifth-title]
port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name. Either the UUID or
both names may be provided on input.

|enabled
|boolean
a|

|lag
|link:#lag[lag]
a|

|mac_address
|string
a|

|mtu

```

```

|integer
a|MTU of the port in bytes. Set by broadcast domain.

|name
|string
a|Portname, such as e0a, e1b-100 (VLAN on ethernet), a0c (LAG/ifgrp), a0d-
200 (vlan on LAG/ifgrp)

|reachability
|string
a|Reachability status of the port. Enum value "ok" is the only acceptable
value for a PATCH request to repair a port.

|reachable_broadcast_domains
|array[link:#reachable_broadcast_domains[reachable_broadcast_domains]]
a|Reachable broadcast domains.

|speed
|integer
a|Link speed in Mbps

|state
|string
a|Operational state of the port.

|uuid
|string
a|Port UUID

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===


```
//end collapsible .Definitions block
====

[[IDd326c985045c927a2e45f5409398e133]]
= Retrieve historical port performance metrics

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/ports/{uuid}/metrics`#

*Introduced In:* 9.8

Retrieves historical performance metrics for a port.

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|timestamp
|string
|query
|False
a|Filter by timestamp

|duration
|string
|query
|False
a|Filter by duration

|throughput.write
|integer
|query
|False
a|Filter by throughput.write
```

```
|throughput.read
|integer
|query
|False
a|Filter by throughput.read
```

```
|throughput.total
|integer
|query
|False
a|Filter by throughput.total
```

```
|status
|string
|query
|False
a|Filter by status
```

```
|uuid
|string
|path
|True
a|Unique identifier of the port.
```

```
|interval
|string
|query
|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:
```

```
* 1h: Metrics over the most recent hour sampled over 15 seconds.
* 1d: Metrics over the most recent day sampled over 5 minutes.
* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]
```

```
|return_timeout
```

```

|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records


|records
|array[link:#records[records]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "status": "ok",

```

```

    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

```

====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the port object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#records]
[.api-collapsible-fifth-title]
records
```

Throughput performance for the Ethernet port.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"inconsistent_old_data" is returned when one or more nodes do not have the
latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the port object.
```

```
|timestamp
|string
```


a|The timestamp of the performance data.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
= Retrieve network Ethernet switch ports
```

```
:leveloffset: +1
```

```
[[IDeafe38ec34beb0f6c9887790f7074aed]]
= Network Ethernet switch ports endpoint overview
```

```
== Overview
```

This API can be used to get the port information for an ethernet switch used in a cluster or storage networks. This API supports GET only. The GET operation returns a list of ports with status and configuration information.

```
== Examples
```

```
=== Retrieving the ports for ethernet switches
```

The following example retrieves the ethernet switch ports for all the ethernet switches used for cluster and/or storage networks. Note that if the `_fields=*` parameter is not specified, the fields in-

octets, in-errors, in-discards, out-octets, out-errors, out-discards, interface-number, unique-name, mac-address are not returned. Filters can be added on the fields to limit the results.

```
[,json]
```

```
----
```

```
# The API:
```

```
GET /network/ethernet/switch/ports
```

```
# The call:
```

```
curl -X GET "https://<mgmt-  
ip>/api/network/ethernet/switch/ports?return_records=true" -H "accept:  
application/json" -H "Content-Type: application/hal+json"
```

```
# The response: (abbreviated output due to length, shows a port connected  
to a node, a port with no remote connection and a port connected to shelf)
```

```
{  
  "records": [  
    {  
      "switch": {  
        "name": "RTP-CS01-510R11(FOC22092K12)",  
        "_links": {  
          "self": {  
            "href": "/api/network/ethernet/switches/RTP-CS01-  
510R11(FOC22092K12)"  
          }  
        }  
      },  
      "identity": {  
        "name": "Ethernet1/1",  
        "index": 436207616,  
        "number": 1,  
        "type": "ethernetcsmacd"  
      },  
      "mtu": 9216,  
      "duplex_type": "full_duplex",  
      "speed": 100000,  
      "configured": "up",  
      "state": "up",  
      "isl": false,  
      "statistics": {  
        "receive_raw": {  
          "packets": 1616467751,  
          "errors": 0,  
          "discards": 0
```

```

    },
    "transmit_raw": {
      "packets": 206717534,
      "errors": 0,
      "discards": 0
    }
  },
  "remote_port": {
    "mtu": 9000,
    "name": "e3a",
    "device": {
      "node": {
        "name": "stiA400-311",
        "uuid": "54c0f036-8a3a-11ea-893d-00a098fd726d",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/54c0f036-8a3a-11ea-893d-00a098fd726d"
          }
        }
      }
    }
  },
  "mac_address": "00be75ae2ad4",
  "vlan_id": [
    1,
    17,
    18,
    92
  ],
  "_links": {
    "self": {
      "href": "/api/network/ethernet/switch/ports/RTP-CS01-510R11%28FOC22092K12%29/Ethernet1%2F1/436207616"
    }
  }
},
{
  "switch": {
    "name": "RTP-CS01-510R11 (FOC22092K12)",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switches/RTP-CS01-510R11 (FOC22092K12)"
      }
    }
  }
}

```

```

    },
    "identity": {
      "name": "Ethernet1/11",
      "index": 436212736,
      "number": 11,
      "type": "ethernetcsmacd"
    },
    "mtu": 9216,
    "duplex_type": "unknown",
    "speed": 100000,
    "configured": "up",
    "state": "down",
    "isl": false,
    "statistics": {
      "receive_raw": {
        "packets": 0,
        "errors": 0,
        "discards": 0
      },
      "transmit_raw": {
        "packets": 0,
        "errors": 0,
        "discards": 0
      }
    },
    "mac_address": "00be75ae2afc",
    "vlan_id": [
      1,
      17,
      18,
      92
    ],
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switch/ports/RTP-CS01-510R11%28FOC22092K12%29/Ethernet1%2F11/436212736"
      }
    }
  },
  {
    "switch": {
      "name": "RTP-SS01-510R10 (FOC22170DFR) ",
      "_links": {
        "self": {
          "href": "/api/network/ethernet/switches/RTP-SS01-510R10 (FOC22170DFR) "
        }
      }
    }
  }
}

```

```

    }
  },
  "identity": {
    "name": "Ethernet1/10",
    "index": 436212224,
    "number": 10,
    "type": "ethernetcsmacd"
  },
  "mtu": 9216,
  "duplex_type": "full_duplex",
  "speed": 100000,
  "configured": "up",
  "state": "up",
  "isl": false,
  "statistics": {
    "receive_raw": {
      "packets": 332013844,
      "errors": 0,
      "discards": 0
    },
    "transmit_raw": {
      "packets": 2429595607,
      "errors": 0,
      "discards": 0
    }
  },
  "remote_port": {
    "mtu": 9000,
    "name": "e0a",
    "device": {
      "shelf": {
        "name": "SHFFG1828000004:B",
        "uid": "12439000444923584512",
        "_links": {
          "self": {
            "href": "/api/storage/shelves/12439000444923584512"
          }
        }
      }
    }
  },
  "mac_address": "00fcbaead548",
  "vlan_id": [
    1,
    30
  ]
}

```

```

    ],
    "_links": {
        "self": {
            "href": "/api/network/ethernet/switch/ports/RTP-SS01-
510R10%28FOC22170DFR%29/Ethernet1%2F10/436212224"
        }
    }
}
],
"num_records": 138,
"_links": {
    "self": {
        "href":
"/api/network/ethernet/switch/ports?fields=*&return_records=true"
    }
}
}
----

'''

=== Retrieving a ports on an ethernet switch

[,json]
----

# The API:
GET
/network/ethernet/switch/ports/{switch.name}/{identity.name}/{identity.ind
ex}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ethernet/switch/ports/RTP-SS02-
510R10%28FOC22131U6T%29/Ethernet1%2F9/436211712?return_records=true" -H
"accept: application/json" -H "Content-Type: application/hal+json"

# The response:
{
"switch": {
    "name": "RTP-SS02-510R10 (FOC22131U6T) ",
    "_links": {
        "self": {
            "href": "/api/network/ethernet/switches/RTP-SS02-
510R10 (FOC22131U6T) "
        }
    }
}
}

```

```

},
"identity": {
  "name": "Ethernet1/9",
  "index": 436211712,
  "number": 9,
  "type": "ethernetcsmaacd"
},
"mtu": 9216,
"duplex_type": "full_duplex",
"speed": 100000,
"configured": "up",
"state": "up",
"isl": false,
"statistics": {
  "receive_raw": {
    "packets": 4012559315,
    "errors": 0,
    "discards": 0
  },
  "transmit_raw": {
    "packets": 337898026,
    "errors": 0,
    "discards": 0
  }
},
"remote_port": {
  "mtu": 9000,
  "name": "e0b",
  "device": {
    "shelf": {
      "name": "SHFFG1828000004:A",
      "uid": "12439000444923584512",
      "_links": {
        "self": {
          "href": "/api/storage/shelves/12439000444923584512"
        }
      }
    }
  }
},
"mac_address": "00fcbaea7228",
"vlan_id": [
  1,
  30
],
"_links": {

```



```

"self": {
  "href": "/api/network/ethernet/switch/ports/RTP-SS02-
510R10%28FOC22131U6T%29/Ethernet1%2F9/436211712"
}
}
}
----

```

```
'''
```

```

[[ID021fb1cd01894fb56018291d4dda9668]]
= Retrieve Ethernet switch ports

```

```

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/switch/ports`#

```

Introduced In: 9.8

Retrieves the ethernet switch ports.

== Related ONTAP commands

* `system switch ethernet interface show`

== Learn more

*

xref:{relative_path}network_ethernet_switch_ports_endpoint_overview.html[D
OC /network/ethernet/switch/ports]

== Parameters

```

[cols=5*,options=header]
|===

```

```

|Name
|Type
|In
|Required
|Description

|state

```

```
|string
|query
|False
a|Filter by state
```

```
|duplex_type
|string
|query
|False
a|Filter by duplex_type
```

```
|identity.name
|string
|query
|False
a|Filter by identity.name
```

```
|identity.index
|integer
|query
|False
a|Filter by identity.index
```

```
|identity.number
|integer
|query
|False
a|Filter by identity.number
```

```
|switch.name
|string
|query
|False
a|Filter by switch.name
```

```
|statistics.receive_raw.packets
|integer
|query
|False
a|Filter by statistics.receive_raw.packets
```

```
|statistics.receive_raw.discards
|integer
|query
|False
a|Filter by statistics.receive_raw.discards
```

```
|statistics.receive_raw.errors
|integer
|query
|False
a|Filter by statistics.receive_raw.errors
```

```
|statistics.transmit_raw.packets
|integer
|query
|False
a|Filter by statistics.transmit_raw.packets
```

```
|statistics.transmit_raw.discards
|integer
|query
|False
a|Filter by statistics.transmit_raw.discards
```

```
|statistics.transmit_raw.errors
|integer
|query
|False
a|Filter by statistics.transmit_raw.errors
```

```
|configured
|string
|query
|False
a|Filter by configured
```

```
|mtu
|integer
|query
|False
```

a|Filter by mtu

|isl

|boolean

|query

|False

a|Filter by isl

|vlan_id

|integer

|query

|False

a|Filter by vlan_id

|mac_address

|string

|query

|False

a|Filter by mac_address

|type

|string

|query

|False

a|Filter by type

|speed

|integer

|query

|False

a|Filter by speed

|remote_port.mtu

|integer

|query

|False

a|Filter by remote_port.mtu

|remote_port.name

|string

```

|query
|False
a|Filter by remote_port.name

|remote_port.device.shelf.uid
|string
|query
|False
a|Filter by remote_port.device.shelf.uid

|remote_port.device.node.name
|string
|query
|False
a|Filter by remote_port.device.node.name

|remote_port.device.node.uuid
|string
|query
|False
a|Filter by remote_port.device.node.uuid

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

```

* Default value: 1

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#collection_links[collection_links]

a|

|num_records

|integer

a|Number of Records

|records

```
|array[link:#switch_port[switch_port]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "configured": "string",
      "duplex_type": "string",
      "identity": {
        "index": 0,
        "name": "string",
        "number": 0
      },
      "mac_address": "string",
      "mtu": 0,
      "remote_port": {
        "device": {
          "node": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "node1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },

```

```

        "shelf": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "uid": "7777841915827391056"
        },
        "mtu": 0,
        "name": "string"
    },
    "speed": 0,
    "state": "string",
    "statistics": {
        "receive_raw": {
            "discards": "100",
            "errors": "200",
            "packets": "500"
        },
        "transmit_raw": {
            "discards": "100",
            "errors": "200",
            "packets": "500"
        }
    },
    "switch": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "RTP-SS01-510R03 (FOC223443KQ) "
    },
    "type": "string",
    "vlan_id": [
        "integer"
    ]
}
]
}
====

== Error

```

Status: Default, Error


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|href
|string
a|

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

```

```
[#identity]
[.api-collapsible-fifth-title]
identity
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|index
|integer
a|Interface Index.
```

```
|name
|string
a|Interface Name.
```

```
|number
|integer
a|Interface Number.
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#node]
```

```
[.api-collapsible-fifth-title]
```

```
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#shelf]
```

```
[.api-collapsible-fifth-title]
```

```
shelf
```

```
Shelf connected to this port.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|uid
```

```
|string
```

```
a|
```

```
|===
```

```

[#device]
[.api-collapsible-fifth-title]
device

Device connected to port.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|node
|link:#node[node]
a|

|shelf
|link:#shelf[shelf]
a|Shelf connected to this port.


|===


[#remote_port]
[.api-collapsible-fifth-title]
remote_port

Remote port


[cols=3*,options=header]
|===
|Name
|Type
|Description

|device
|link:#device[device]
a|Device connected to port.


|mtu
|integer
a|MTU in octets

```

```
|name
|string
a|Port Name.
```

```
|===
```

```
[#receive_raw]
[.api-collapsible-fifth-title]
receive_raw
```

Packet receive counters for the Ethernet port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|discards
|integer
```

```
a|Total number of discarded packets.
```

```
|errors
|integer
```

```
a|Number of packet errors.
```

```
|packets
|integer
```

```
a|Total packet count.
```

```
|===
```

```
[#transmit_raw]
[.api-collapsible-fifth-title]
transmit_raw
```

Packet transmit counters for the Ethernet port.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|discards
|integer
a|Total number of discarded packets.
```

```
|errors
|integer
a|Number of packet errors.
```

```
|packets
|integer
a|Total packet count.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw counters for the device associated with the Ethernet port.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|receive_raw
|link:#receive_raw[receive_raw]
a|Packet receive counters for the Ethernet port.
```

```
|transmit_raw
|link:#transmit_raw[transmit_raw]
a|Packet transmit counters for the Ethernet port.
```

```
|===
```

```
[#switch]  
[.api-collapsible-fifth-title]  
switch
```

The name of the specified cluster or storage switch.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description  
  
|_links  
|link:#self_link[self_link]  
a|
```

```
|name  
|string  
a|
```

```
|===
```

```
[#switch_port]  
[.api-collapsible-fifth-title]  
switch_port
```

Ethernet Switch Port REST API

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description  
  
|_links  
|link:#self_link[self_link]  
a|
```

```
|configured  
|string
```


a|Administrative Status.

|duplex_type

|string

a|Duplex Settings.

|identity

|link:#identity[identity]

a|

|isl

|boolean

a|Is configured as an ISL link.

|mac_address

|string

a|MAC Address.

|mtu

|integer

a|MTU.

|remote_port

|link:#remote_port[remote_port]

a|Remote port

|speed

|integer

a|Interface Speed(Mbps)

|state

|string

a|Operational Status.

|statistics

|link:#statistics[statistics]

a|These are raw counters for the device associated with the Ethernet port.

```
|switch
|link:#switch[switch]
a|The name of the specified cluster or storage switch.
```

```
|type
|string
a|Interface Type.
```

```
|vlan_id
|array[integer]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
```

```

|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID1fe3b900a79801f37ce69a92567b63f7]]
= Retrieve an Ethernet switch port

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/switch/ports/{switch}/{identity.name}/{identity.
index}`#

*Introduced In:* 9.8

Retrieves an ethernet switch port.

== Related ONTAP commands

* `system switch ethernet interface show`

```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|switch
|string
|path
|True
a|Switch Name
```

```
|identity.name
|string
|path
|True
a|Interface Name
```

```
|identity.index
|integer
|path
|True
a|Interface Index
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
```

```

|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|configured
|string
a|Administrative Status.

|duplex_type
|string
a|Duplex Settings.

|identity
|link:#identity[identity]
a|

|isl
|boolean
a|Is configured as an ISL link.

|mac_address
|string
a|MAC Address.

|mtu
|integer
a|MTU.

|remote_port
|link:#remote_port[remote_port]
a|Remote port

|speed
|integer
a|Interface Speed(Mbps)

```

```
|state
|string
a|Operational Status.
```

```
|statistics
|link:#statistics[statistics]
a|These are raw counters for the device associated with the Ethernet port.
```

```
|switch
|link:#switch[switch]
a|The name of the specified cluster or storage switch.
```

```
|type
|string
a|Interface Type.
```

```
|vlan_id
|array[integer]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "configured": "string",
  "duplex_type": "string",
  "identity": {
    "index": 0,
    "name": "string",
    "number": 0
  },
  "mac_address": "string",
  "mtu": 0,
  "remote_port": {
```

```

"device": {
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
  },
  "shelf": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uid": "7777841915827391056"
  }
},
"mtu": 0,
"name": "string"
},
"speed": 0,
"state": "string",
"statistics": {
  "receive_raw": {
    "discards": "100",
    "errors": "200",
    "packets": "500"
  },
  "transmit_raw": {
    "discards": "100",
    "errors": "200",
    "packets": "500"
  }
},
"switch": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "RTP-SS01-510R03(FOC223443KQ) "
},
"type": "string",
"vlan_id": [

```

```

    "integer"
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]

```



```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#self_link]
```

```
[.api-collapsible-fifth-title]
```

```
self_link
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#identity]
```

```
[.api-collapsible-fifth-title]
```

```
identity
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|index
```

```

|integer
a|Interface Index.

|name
|string
a|Interface Name.

|number
|integer
a|Interface Number.

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|name  
|string  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#shelf]  
[.api-collapsible-fifth-title]  
shelf
```

Shelf connected to this port.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|uid  
|string  
a|
```

```
|===
```

```
[#device]  
[.api-collapsible-fifth-title]  
device
```

Device connected to port.

```
[cols=3*,options=header]  
|===  
|Name  
|Type
```

```

|Description

|node
|link:#node[node]
a|

|shelf
|link:#shelf[shelf]
a|Shelf connected to this port.

|===

[#remote_port]
[.api-collapsible-fifth-title]
remote_port

Remote port

[cols=3*,options=header]
|===
|Name
|Type
|Description

|device
|link:#device[device]
a|Device connected to port.

|mtu
|integer
a|MTU in octets

|name
|string
a|Port Name.

|===

[#receive_raw]
[.api-collapsible-fifth-title]

```

```
receive_raw
```

Packet receive counters for the Ethernet port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|discards
```

```
|integer
```

```
a|Total number of discarded packets.
```

```
|errors
```

```
|integer
```

```
a|Number of packet errors.
```

```
|packets
```

```
|integer
```

```
a|Total packet count.
```

```
|===
```

```
[#transmit_raw]
```

```
[.api-collapsible-fifth-title]
```

```
transmit_raw
```

Packet transmit counters for the Ethernet port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|discards
```

```
|integer
```

```
a|Total number of discarded packets.
```

```
|errors
|integer
a|Number of packet errors.
```

```
|packets
|integer
a|Total packet count.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw counters for the device associated with the Ethernet port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|receive_raw
|link:#receive_raw[receive_raw]
a|Packet receive counters for the Ethernet port.
```

```
|transmit_raw
|link:#transmit_raw[transmit_raw]
a|Packet transmit counters for the Ethernet port.
```

```
|===
```

```
[#switch]
[.api-collapsible-fifth-title]
switch
```

The name of the specified cluster or storage switch.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|name
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage Ethernet switches

:leveloffset: +1

[[ID6ad14c4bb6c8e0d79f574262787b61d2]]
= Network Ethernet switches endpoint overview


== Overview

```

This API can be used to get information about the ethernet switches used for cluster and/or storage networks. This API support GET and PATCH calls.

The GET operation returns a list of discovered switches with status and configuration information. The PATCH can be used to modify state of the switch.

== Examples

=== Retrieving the ethernet switches for a cluster

The following example retrieves the ONTAP switches from the cluster. Note that if the `_fields=*` parameter is not specified, the fields `snmp.version`, `snmp.credential`, `model`, `sw-version`, `reason`, `version-source`, `monitoring.enable`, `monitoring.subsystem.type` and `monitoring.subsystem.poll_interval` are not returned. Filters can be added on the fields to limit the results.

```
[,json]
```

```
----
```

```
# The API:
```

```
GET network/ethernet/switches
```

```
# The call:
```

```
curl -X GET "https://<mgmt-  
ip>/api/network/ethernet/switches?return_records=true" -H "accept:  
application/json" -H "Content-Type: application/hal+json"
```

```
# The response:
```

```
{  
  "records": [  
    {  
      "name": "RTP-CS01-510R11(FOC22092K12)",  
      "address": "172.26.207.77",  
      "discovered": true,  
      "model": "NX3232C",  
      "monitoring": {  
        "monitored": true,  
        "reason": "None"  
      },  
      "network": "cluster",  
      "serial_number": "Unknown",  
      "snmp": {  
        "version": "snmpv2c",  
        "user": "cshml!"  
      },  
      "version": "Cisco Nexus Operating System (NX-OS) Software, Version  
9.2(3)",  
    },  
  ],  
}
```

```

    "_links": {
      "self": {
        "href": "/api/network/ethernet/switches/RTP-CS01-510R11%28FOC22092K12%29"
      }
    }
  },
  {
    "name": "RTP-CS01-510R12 (FOC22373C3P)",
    "address": "172.26.207.82",
    "discovered": true,
    "model": "NX3232C",
    "monitoring": {
      "monitored": true,
      "reason": "None"
    },
    "network": "cluster",
    "serial_number": "FOC22373C3P",
    "snmp": {
      "version": "snmpv2c",
      "user": "cshml!"
    },
    "version": "Cisco Nexus Operating System (NX-OS) Software, Version 9.2(3)",
    "_links": {
      "self": {
        "href": "/api/network/ethernet/switches/RTP-CS01-510R12%28FOC22373C3P%29"
      }
    }
  },
  {
    "name": "RTP-SS01-510R10 (FOC22170DFR)",
    "address": "172.26.207.65",
    "discovered": true,
    "model": "NX3232C",
    "monitoring": {
      "monitored": true,
      "reason": "None"
    },
    "network": "storage",
    "serial_number": "FOC22170DFR",
    "snmp": {
      "version": "snmpv2c",
      "user": "cshml!"
    },
    "version": "Cisco Nexus Operating System (NX-OS) Software, Version

```

```

9.3(3)",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/switches/RTP-SS01-
510R10%28FOC22170DFR%29"
    }
  }
},
{
  "name": "RTP-SS02-510R10(FOC22131U6T)",
  "address": "172.26.207.66",
  "discovered": true,
  "model": "NX3232C",
  "monitoring": {
    "monitored": true,
    "reason": "None"
  },
  "network": "storage",
  "serial_number": "FOC22131U6T",
  "snmp": {
    "version": "snmpv2c",
    "user": "cshml!"
  },
  "version": "Cisco Nexus Operating System (NX-OS) Software, Version
9.3(3)",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/switches/RTP-SS02-
510R10%28FOC22131U6T%29"
    }
  }
}
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/network/ethernet/switches?fields=*&return_records=true"
  }
}
}
}
----

'''

```

=== Retrieving an ethernet switch for a cluster

The following example retrieves a single switch by the switchname using the API.

```
[,json]
```

```
----
```

```
# The API:
```

```
GET /network/ethernet/switches{name}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/network/ethernet/switches/RTP-SS02-510R10(FOC22131U6T?return_records=true" -H "accept: application/json" -H "Content-Type: application/hal+json"
```

```
# The response:
```

```
{
  "name": "RTP-SS02-510R10(FOC22131U6T)",
  "address": "172.26.207.66",
  "discovered": true,
  "model": "NX3232C",
  "monitoring": {
    "monitored": true,
    "reason": "None"
  },
  "network": "storage",
  "serial_number": "FOC22131U6T",
  "snmp": {
    "version": "snmpv2c",
    "user": "cshml!"
  },
  "version": "Cisco Nexus Operating System (NX-OS) Software, Version 9.3(3)",
  "_links": {
    "self": {
      "href": "/api/network/ethernet/switches/RTP-SS02-510R10(FOC22131U6T)"
    }
  }
}
```

```
----
```

```
'''
```

```
=== Configuring a switch
```

The following example configures SNMP credential and version on a switch.

```
[,json]
----

# The API:
PATCH /network/ethernet/switches{name}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ethernet/switches/sconga-
corduroyl-03" -H "accept: application/json" -H "Content-Type:
application/hal+json" -d "{ \"snmp\": { \"version\": \"snmpv2c\",
\"user\": \"cshml!\"} }"

# The response:
{}
----

'''

[[IDf6615329caa2e08f6e283e0fa9a2b6e7]]
= Retrieve Ethernet switches attached to a chassis

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/switches`#

*Introduced In:* 9.8

Retrieves the ethernet switches attached to the chassis.

== Related ONTAP commands

* `system switch ethernet show`

== Learn more

* xref:{relative_path}network_ethernet_switches_endpoint_overview.html[DOC
/network/ethernet/switches]

== Parameters

[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|serial_number
|string
|query
|False
a|Filter by serial_number
```

```
|address
|string
|query
|False
a|Filter by address
```

```
|discovered
|boolean
|query
|False
a|Filter by discovered
```

```
|snmp.version
|string
|query
|False
a|Filter by snmp.version
```

```
|snmp.user
|string
|query
|False
a|Filter by snmp.user
```

```
|monitoring.monitored
|boolean
|query
|False
a|Filter by monitoring.monitored
```

```
|monitoring.reason
|string
|query
|False
a|Filter by monitoring.reason
```

```
|monitoring.enabled
|boolean
|query
|False
a|Filter by monitoring.enabled
```

```
|network
|string
|query
|False
a|Filter by network
```

```
|model
|string
|query
|False
a|Filter by model
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|version
|string
|query
|False
a|Filter by version
```

```
|fields
|array[string]
|query
|False
```

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|return_records

|boolean

|query

|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of Records

|records
|array[link:#switch[switch]]
a|

|===

```

```

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "address": "string",
      "model": "string",
      "monitoring": {
        "reason": "string"
      }
    }
  ]
}

```

```

    },
    "name": "string",
    "network": "string",
    "serial_number": "string",
    "snmp": {
      "user": "string",
      "version": "string"
    },
    "version": "string"
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",

```

```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

```

```

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#monitoring]
[.api-collapsible-fifth-title]
monitoring

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Enable Health Monitoring.

|monitored
|boolean
a|Is Monitored.

|reason
|string
a|Reason For Not Monitoring.

|===

```

```

[#snmp]
[.api-collapsible-fifth-title]
snmp

[cols=3*,options=header]
|===
|Name
|Type
|Description

|user
|string
a|Community String or SNMPv3 Username.

|version
|string
a|SNMP Version.

|===

[#switch]
[.api-collapsible-fifth-title]
switch

Ethernet Switch REST API

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|address
|string
a|IP Address.

|discovered

```

```

|boolean
a|Discovered By ONTAP CDP/LLDP

* readOnly: 1
* Introduced in: 9.8


|model
|string
a|Model Number.


|monitoring
|link:#monitoring[monitoring]
a|

|name
|string
a|Name.


|network
|string
a|Switch Network.


|serial_number
|string
a|Serial Number.


|snmp
|link:#snmp[snmp]
a|

|version
|string
a|Software Version.


|===


[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

|===

//end collapsible .Definitions block
=====

[[IDc695ce6de2b29a723938457880db78d1]]
= Retrieve Ethernet switch details

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ethernet/switches/{name}`#

Introduced In: 9.8

Retrieves the details of an ethernet switch.

== Related ONTAP commands

* `system switch ethernet show`

== Learn more

* xref:{relative_path}network_ethernet_switches_endpoint_overview.html[DOC
/network/ethernet/switches]

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

name
string
path
True
a Name


```
|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|address
|string
a|IP Address.

|discovered
|boolean
a|Discovered By ONTAP CDP/LLDP

* readOnly: 1
* Introduced in: 9.8

|model
|string
a|Model Number.

|monitoring
|link:#monitoring[monitoring]
a|

|name
```

```
|string
```

```
a|Name.
```

```
|network
```

```
|string
```

```
a|Switch Network.
```

```
|serial_number
```

```
|string
```

```
a|Serial Number.
```

```
|snmp
```

```
|link:#snmp[snmp]
```

```
a|
```

```
|version
```

```
|string
```

```
a|Software Version.
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{  
  "_links": {  
    "self": {  
      "href": "/api/resourcelink"  
    }  
  },  
  "address": "string",  
  "model": "string",  
  "monitoring": {  
    "reason": "string"  
  },  
  "name": "string",  
  "network": "string",  
  "serial_number": "string",  
  "snmp": {  
    "user": "string",  
    "version": "string"
```

```

    },
    "version": "string"
  }
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]

```

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#self_link]
```

```
[.api-collapsible-fifth-title]
```

```
self_link
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#monitoring]
```

```
[.api-collapsible-fifth-title]
```

```
monitoring
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```

|boolean
a|Enable Health Monitoring.

|monitored
|boolean
a|Is Monitored.

|reason
|string
a|Reason For Not Monitoring.

|===

[#snmp]
[.api-collapsible-fifth-title]
snmp

[cols=3*,options=header]
|===
|Name
|Type
|Description

|user
|string
a|Community String or SNMPv3 Username.

|version
|string
a|SNMP Version.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument


|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

```

```
//end collapsible .Definitions block
====
```

```
[[IDcb5b9610591e4ba525154544954e156a]]
= Update an Ethernet switch
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/network/ethernet/switches/{name}`#
```

```
*Introduced In:* 9.8
```

```
Update Ethernet Switch REST API
```

```
== Parameters
```

```
[cols=5*,options=header]
|==
```

```
|Name
|Type
|In
|Required
|Description
```

```
|name
|string
|path
|True
a|Switch Name
```

```
|return_timeout
|integer
|query
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1
* Max value: 120
* Min value: 0

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|address

|string

a|IP Address.

|discovered

|boolean

a|Discovered By ONTAP CDP/LLDP

* readOnly: 1

* Introduced in: 9.8

|model

|string

a|Model Number.

|monitoring

|link:#monitoring[monitoring]

a|

|name

|string

a|Name.

|network

|string

a|Switch Network.


```
|serial_number
|string
a|Serial Number.
```

```
|snmp
|link:#snmp[snmp]
a|
```

```
|version
|string
a|Software Version.
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "address": "string",
  "model": "string",
  "monitoring": {
    "reason": "string"
  },
  "name": "string",
  "network": "string",
  "serial_number": "string",
  "snmp": {
    "user": "string",
    "version": "string"
  },
  "version": "string"
}
====
```

```
== Response
```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Error

```

Status: Default, Error

```

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link
[#monitoring]
[.api-collapsible-fifth-title]
monitoring

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Enable Health Monitoring.

|monitored
|boolean
a|Is Monitored.

|reason
|string
a|Reason For Not Monitoring.

|===

[#snmp]
[.api-collapsible-fifth-title]
snmp

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|user
|string
a|Community String or SNMPv3 Username.
```

```
|version
|string
a|SNMP Version.
```

```
|===
```

```
[#switch]
[.api-collapsible-fifth-title]
switch
```

Ethernet Switch REST API

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|address
|string
a|IP Address.
```

```
|discovered
|boolean
a|Discovered By ONTAP CDP/LLDP
```

```
* readOnly: 1
* Introduced in: 9.8
```

```
|model
|string
a|Model Number.
```

```
|monitoring
|link:#monitoring[monitoring]
```

```

a|

|name
|string
a|Name.

|network
|string
a|Switch Network.

|serial_number
|string
a|Serial Number.

|snmp
|link:#snmp[snmp]
a|

|version
|string
a|Software Version.

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

```

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

|===

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```

|===

```

```

//end collapsible .Definitions block
====

```

```

:leveloffset: -1

```

```

= Manage FC network interfaces

```

```

:leveloffset: +1

```

```

[[ID00b11ca1b528b4f5fb1b3b2df2d0ce0d]]
= Network FC interfaces endpoint overview

```

```

== Overview

```

Fibre Channel (FC) interfaces are the logical endpoints for FC network connections to an SVM. An FC interface provides FC access to storage within the interface SVM using either Fibre Channel Protocol or NVMe over FC (NVMe/FC).

The Fibre Channel interface REST API allows you to create, delete, update, and discover FC interfaces, and obtain status information for FC interfaces.

An FC interface is created on an FC port which is located on a cluster node. The FC port must be specified to identify the location of the interface for a POST or PATCH request that relocates an interface. You can identify the port by supplying either the node and port names or the port UUID.

== Performance monitoring

Performance of an FC interface can be monitored by observing the ``metric.++`` and ``statistics.++`` properties. These properties show the performance of an FC interface in terms of IOPS, latency, and throughput. The ``metric.++`` properties denote an average, whereas ``statistics.++`` properties denote a real-time monotonically increasing value aggregated across all nodes.

== Examples

=== Creating an FC interface using the port node and name to identify the location

This example uses the ``return_records`` query parameter to retrieve the newly created FC interface in the POST response.

The API:

POST /api/network/fc/interfaces

The call:

```
curl -X POST 'https://<mgmt-  
ip>/api/network/fc/interfaces?return_records=true' -H 'accept:  
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name": "lif1",  
"location": { "home_port": { "name": "0a", "home_node": { "name": "node1"  
} } }, "data_protocol": "fc" }'
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "svm": {  
        "uuid": "cf300f5c-db83-11e8-bd46-005056bba0e0",  
        "name": "svm1",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/cf300f5c-db83-11e8-bd46-005056bba0e0"  
          }  
        }  
      },  
      "uuid": "f6045b92-dec7-11e8-a733-005056bba0e0",  
      "name": "lif1",
```



```

"location": {
  "home_node": {
    "uuid": "bafe9b9f-db81-11e8-bd46-005056bba0e0",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/bafe9b9f-db81-11e8-bd46-005056bba0e0"
      }
    }
  },
  "home_port": {
    "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
    "name": "0a",
    "node": {
      "name": "node1"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-005056bba0e0"
      }
    }
  },
  "node": {
    "uuid": "bafe9b9f-db81-11e8-bd46-005056bba0e0",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/bafe9b9f-db81-11e8-bd46-005056bba0e0"
      }
    }
  },
  "port": {
    "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
    "name": "0a",
    "node": {
      "name": "node1"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-005056bba0e0"
      }
    }
  }
}

```

```

    }
  },
  "enabled": true,
  "state": "down",
  "data_protocol": "fc",
  "wwpn": "20:04:00:50:56:bb:a0:e0",
  "wwnn": "20:00:00:50:56:bb:a0:e0",
  "port_address": "9da2cb1",
  "_links": {
    "self": {
      "href": "/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-005056bba0e0"
    }
  }
}
]
}

```

'''

=== Creating an FC interface using the port UUID to identify the location

This example uses the `return_records` query parameter to retrieve the newly created FC interface in the POST response.

The API:

POST /api/network/fc/interfaces

The call:

```

curl -X POST 'https://<mgmt-
ip>/api/network/fc/interfaces?return_records=true' -H 'accept:
application/hal+json' -d '{ "svm": { "name": "svm3" }, "name": "lif2",
"location": { "home_port": { "uuid": "24bb636a-db83-11e8-9a49-
005056bb1ec6" } }, "data_protocol": "fc_nvme" }'

```

The response:

```

{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "uuid": "a5060466-dbab-11e8-bd46-005056bba0e0",
        "name": "svm3",

```

```

    "_links": {
      "self": {
        "href": "/api/svm/svms/a5060466-dbab-11e8-bd46-005056bba0e0"
      }
    },
    "uuid": "cdeb5591-dec9-11e8-a733-005056bba0e0",
    "name": "lif2",
    "location": {
      "home_node": {
        "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
        "name": "node3",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-
005056bb1ec6"
          }
        }
      },
      "home_port": {
        "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
        "name": "1b",
        "node": {
          "name": "node3"
        },
        "_links": {
          "self": {
            "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-
005056bb1ec6"
          }
        }
      },
      "node": {
        "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
        "name": "node3",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-
005056bb1ec6"
          }
        }
      },
      "port": {
        "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
        "name": "1b",
        "node": {

```

```

        "name": "node3"
    },
    "_links": {
        "self": {
            "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-005056bb1ec6"
        }
    }
}
},
"enabled": true,
"state": "down",
"data_protocol": "fc_nvme",
"wwpn": "20:05:00:50:56:bb:a0:e0",
"wwnn": "20:02:00:50:56:bb:a0:e0",
"port_address": "612e202b",
"_links": {
    "self": {
        "href": "/api/network/fc/interfaces/cdeb5591-dec9-11e8-a733-005056bba0e0"
    }
}
}
]
}
----

'''

```

=== Retrieving all properties for all FC interfaces

This example uses the `fields` query parameter to retrieve all properties.

The API:

```
GET /api/network/fc/interfaces
```

The call:

```
curl -X GET 'https://<mgmt-ip>/api/network/fc/interfaces?fields=*' -H
'accept: application/hal+json'
```

The response:

```
{
"records": [
{

```

```

"svm": {
  "uuid": "a5060466-dbab-11e8-bd46-005056bba0e0",
  "name": "svm3",
  "_links": {
    "self": {
      "href": "/api/svm/svms/a5060466-dbab-11e8-bd46-005056bba0e0"
    }
  }
},
"uuid": "cdeb5591-dec9-11e8-a733-005056bba0e0",
"name": "lif2",
"location": {
  "home_node": {
    "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
    "name": "node3",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-
005056bb1ec6"
      }
    }
  },
  "home_port": {
    "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
    "name": "1b",
    "node": {
      "name": "node3"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-
005056bb1ec6"
      }
    }
  },
  "node": {
    "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
    "name": "node3",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-
005056bb1ec6"
      }
    }
  },
  "port": {

```

```

    "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
    "name": "1b",
    "node": {
      "name": "node3"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-005056bb1ec6"
      }
    }
  },
  "enabled": true,
  "state": "down",
  "data_protocol": "fc_nvme",
  "wwpn": "20:05:00:50:56:bb:a0:e0",
  "wwnn": "20:02:00:50:56:bb:a0:e0",
  "port_address": "612e202b",
  "_links": {
    "self": {
      "href": "/api/network/fc/interfaces/cdeb5591-dec9-11e8-a733-005056bba0e0"
    }
  }
},
{
  "svm": {
    "uuid": "cf300f5c-db83-11e8-bd46-005056bba0e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/cf300f5c-db83-11e8-bd46-005056bba0e0"
      }
    }
  },
  "uuid": "f6045b92-dec7-11e8-a733-005056bba0e0",
  "name": "lif1",
  "location": {
    "home_node": {
      "uuid": "bafe9b9f-db81-11e8-bd46-005056bba0e0",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/bafe9b9f-db81-11e8-bd46-005056bba0e0"
        }
      }
    }
  }
}

```

```

    }
  },
  "home_port": {
    "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
    "name": "0a",
    "node": {
      "name": "node1"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-005056bba0e0"
      }
    }
  },
  "node": {
    "uuid": "baf9b9f-db81-11e8-bd46-005056bba0e0",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/baf9b9f-db81-11e8-bd46-005056bba0e0"
      }
    }
  },
  "port": {
    "uuid": "300c1ae3-db82-11e8-bd46-005056bba0e0",
    "name": "0a",
    "node": {
      "name": "node1"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/300c1ae3-db82-11e8-bd46-005056bba0e0"
      }
    }
  },
  "enabled": true,
  "state": "down",
  "data_protocol": "fcp",
  "wwpn": "20:04:00:50:56:bb:a0:e0",
  "wwnn": "20:00:00:50:56:bb:a0:e0",
  "port_address": "9da2cb1",

```

```

    "_links": {
      "self": {
        "href": "/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-005056bba0e0"
      }
    }
  },
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/network/fc/interfaces?fields=*"
    }
  }
}
----

'''

```

=== Retrieving a list of selected FC interfaces

This example uses property query parameters to retrieve FC interfaces configured for the FC Protocol that are set to `_up_`.

The API:

GET /api/network/fc/interfaces

The call:

```

curl -X GET 'https://<mgmt-ip>/api/network/fc/interfaces?data_protocol=fcp&state=up' -H 'accept: application/hal+json'

```

The response:

```

{
  "records": [
    {
      "svm": {
        "uuid": "cf300f5c-db83-11e8-bd46-005056bba0e0",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/cf300f5c-db83-11e8-bd46-005056bba0e0"
          }
        }
      }
    }
  ]
}

```



```

    },
    "uuid": "f6045b92-dec7-11e8-a733-005056bba0e0",
    "name": "lif1",
    "state": "up",
    "data_protocol": "fcp",
    "_links": {
      "self": {
        "href": "/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-
005056bba0e0"
      }
    }
  },
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/network/fc/interfaces?data_protocol=fcp&state=up"
    }
  }
}
}
}
-----
'''

```

=== Retrieving a specific FC interface

```

-----

# The API:
GET /api/network/fc/interfaces/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/network/fc/interfaces/cdeb5591-dec9-
11e8-a733-005056bba0e0' -H 'accept: application/hal+json'

# The response:
{
  "svm": {
    "uuid": "a5060466-dbab-11e8-bd46-005056bba0e0",
    "name": "svm3",
    "_links": {
      "self": {
        "href": "/api/svm/svms/a5060466-dbab-11e8-bd46-005056bba0e0"
      }
    }
  }
},

```

```

"uuid": "cdeb5591-dec9-11e8-a733-005056bba0e0",
"name": "lif2",
"location": {
  "home_node": {
    "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
    "name": "node3",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-005056bb1ec6"
      }
    }
  },
  "home_port": {
    "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
    "name": "1b",
    "node": {
      "name": "node3"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-005056bb1ec6"
      }
    }
  },
  "node": {
    "uuid": "e85aa147-db83-11e8-9a48-005056bb1ec6",
    "name": "node3",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e85aa147-db83-11e8-9a48-005056bb1ec6"
      }
    }
  },
  "port": {
    "uuid": "24bb636a-db83-11e8-9a49-005056bb1ec6",
    "name": "1b",
    "node": {
      "name": "node3"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/24bb636a-db83-11e8-9a49-005056bb1ec6"
      }
    }
  }
}

```

```

    }
  },
  "enabled": true,
  "state": "down",
  "data_protocol": "fc_nvme",
  "wwpn": "20:05:00:50:56:bb:a0:e0",
  "wwnn": "20:02:00:50:56:bb:a0:e0",
  "port_address": "612e202b",
  "metric": {
    "timestamp": "2019-04-09T05:50:15Z",
    "duration": "PT15S",
    "status": "ok",
    "latency": {
      "other": 0,
      "total": 0,
      "read": 0,
      "write": 0
    },
    "iops": {
      "read": 0,
      "write": 0,
      "other": 0,
      "total": 0
    },
    "throughput": {
      "read": 0,
      "write": 0,
      "total": 0
    }
  },
  "statistics": {
    "timestamp": "2019-04-09T05:50:42Z",
    "status": "ok",
    "latency_raw": {
      "other": 38298,
      "total": 38298,
      "read": 0,
      "write": 0
    },
    "iops_raw": {
      "read": 0,
      "write": 0,
      "other": 3,
      "total": 3
    },
    "throughput_raw": {

```

```

        "read": 0,
        "write": 0,
        "total": 0
    }
},
"_links": {
    "self": {
        "href": "/api/network/fc/interfaces/cdeb5591-dec9-11e8-a733-005056bba0e0"
    }
}
}
}
-----
'''

```

== Disabling an FC interface

When updating certain properties or deleting an FC interface, the interface must first be disabled using the following:

```

-----

# The API:
PATCH /api/network/fc/interfaces/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-005056bba0e0' -H 'accept: application/hal+json' -d '{ "enabled": false }'
-----

'''

```

=== Moving an FC interface to a new node and port

To move an FC interface to another node or port, the destination FC port must be specified in a PATCH request. Either the port UUID or node and port names can be used to identify the port.

Note that only FC interfaces configured for the FC Protocol can be moved. FC interfaces configured for NVMe/FC cannot be moved. The interface must also be set to the disabled state before being moved.

```

-----

# The API:

```

```
PATCH /api/network/fc/interfaces/{uuid}
```

```
# The call:
```

```
curl -X PATCH 'https://<mgmt-ip>/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-005056bba0e0' -H 'accept: application/hal+json' -d '{
  "location": { "home_port": { "uuid": "a1dc7aa5-db83-11e8-9ef7-005056bbbbcc" } } }'
```

```
----
```

```
'''
```

```
=== Deleting an FC interface
```

The FC interface must be disabled before being deleted.

```
----
```

```
# The API:
```

```
DELETE /api/network/fc/interfaces/{uuid}
```

```
# The call:
```

```
curl -X DELETE 'https://<mgmt-ip>/api/network/fc/interfaces/f6045b92-dec7-11e8-a733-005056bba0e0' -H 'accept: application/hal+json'
```

```
----
```

```
[[IDca3d89284f47c299b45b2f4c7b2877ac]]
```

```
= Retrieve FC interfaces
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/network/fc/interfaces`#
```

```
*Introduced In:* 9.6
```

```
Retrieves FC interfaces.
```

```
== Related ONTAP commands
```

```
* `network interface show`
```

```
* `vserver fcp interface show`
```

```
== Learn more
```

```
* xref:{relative_path}network_fc_interfaces_endpoint_overview.html[DOC
```

```
/network/fc/interfaces]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|location.port.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by location.port.name
```

```
|location.port.node.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by location.port.node.name
```

```
|location.port.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by location.port.uuid
```

```
|location.home_node.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by location.home_node.name
```

```
* Introduced in: 9.8
```

```
|location.home_node.uuid
```

```
|string
```

```
|query
```

```
|False
a|Filter by location.home_node.uuid

* Introduced in: 9.8

|location.home_port.name
|string
|query
|False
a|Filter by location.home_port.name

* Introduced in: 9.8

|location.home_port.node.name
|string
|query
|False
a|Filter by location.home_port.node.name

* Introduced in: 9.8

|location.home_port.uuid
|string
|query
|False
a|Filter by location.home_port.uuid

* Introduced in: 9.8

|location.node.name
|string
|query
|False
a|Filter by location.node.name

|location.node.uuid
|string
|query
|False
a|Filter by location.node.uuid
```

```
|location.is_home
|boolean
|query
|False
a|Filter by location.is_home

* Introduced in: 9.8

|enabled
|boolean
|query
|False
a|Filter by enabled

|port_address
|string
|query
|False
a|Filter by port_address

|uuid
|string
|query
|False
a|Filter by uuid

|data_protocol
|string
|query
|False
a|Filter by data_protocol

|metric.latency.total
|integer
|query
|False
a|Filter by metric.latency.total

* Introduced in: 9.8

|metric.latency.read
```



```
|integer
|query
|False
a|Filter by metric.latency.read
```

* Introduced in: 9.8

```
|metric.latency.other
|integer
|query
|False
a|Filter by metric.latency.other
```

* Introduced in: 9.8

```
|metric.latency.write
|integer
|query
|False
a|Filter by metric.latency.write
```

* Introduced in: 9.8

```
|metric.status
|string
|query
|False
a|Filter by metric.status
```

* Introduced in: 9.8

```
|metric.iops.total
|integer
|query
|False
a|Filter by metric.iops.total
```

* Introduced in: 9.8

```
|metric.iops.read
|integer
|query
```

```
|False
a|Filter by metric.iops.read

* Introduced in: 9.8

|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other

* Introduced in: 9.8

|metric.iops.write
|integer
|query
|False
a|Filter by metric.iops.write

* Introduced in: 9.8

|metric.duration
|string
|query
|False
a|Filter by metric.duration

* Introduced in: 9.8

|metric.timestamp
|string
|query
|False
a|Filter by metric.timestamp

* Introduced in: 9.8

|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write
```

* Introduced in: 9.8

```
|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read
```

* Introduced in: 9.8

```
|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total
```

* Introduced in: 9.8

```
|name
|string
|query
|False
a|Filter by name
```

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|statistics.iops_raw.total
|integer
|query
|False
a|Filter by statistics.iops_raw.total
```

* Introduced in: 9.8

```
|statistics.iops_raw.read
|integer
```

```
|query
|False
a|Filter by statistics.iops_raw.read

* Introduced in: 9.8

|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other

* Introduced in: 9.8

|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write

* Introduced in: 9.8

|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total

* Introduced in: 9.8

|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read

* Introduced in: 9.8

|statistics.latency_raw.other
|integer
|query
|False
```

```
a|Filter by statistics.latency_raw.other
```

```
* Introduced in: 9.8
```

```
|statistics.latency_raw.write
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by statistics.latency_raw.write
```

```
* Introduced in: 9.8
```

```
|statistics.timestamp
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by statistics.timestamp
```

```
* Introduced in: 9.8
```

```
|statistics.status
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by statistics.status
```

```
* Introduced in: 9.8
```

```
|statistics.throughput_raw.write
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by statistics.throughput_raw.write
```

```
* Introduced in: 9.8
```

```
|statistics.throughput_raw.read
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by statistics.throughput_raw.read
```

* Introduced in: 9.8

```
|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total
```

* Introduced in: 9.8

```
|wwpn
|string
|query
|False
a|Filter by wwpn
```

```
|wwnn
|string
|query
|False
a|Filter by wwnn
```

```
|state
|string
|query
|False
a|Filter by state
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#fc_interface[fc_interface]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    },
  ],
}
```



```

"comment": "string",
"data_protocol": "string",
"location": {
  "home_node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "home_port": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "0a",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "port": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "0a",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},

```

```

"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"name": "lif1",
"port_address": "5060F",
"state": "string",
"statistics": {
  "iops_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
}

```

```

    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "wwnn": "20:00:00:56:b4:13:01",
    "wwpn": "20:00:00:56:b4:13:a8"
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ]
  }
}

```

```

    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next
|link:#href[href]
a|

```

```

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#home_node]
[.api-collapsible-fifth-title]
home_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

```

```
[#node]
[.api-collapsible-fifth-title]
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#fc_port_reference]
[.api-collapsible-fifth-title]
fc_port_reference
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the FC port.
```

```
|node
```

```
|link:#node[node]
a|The node on which the FC port is located.
```

```
|uuid
|string
a|The unique identifier of the FC port.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its node name and port name. Either the UUID or the node name and port name are required for POST. To move an interface, supply either the UUID or the node name and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|home_node
```

```
|link:#home_node[home_node]
```

```
a|
```

```
|home_port
```

```
|link:#fc_port_reference[fc_port_reference]
```

```
a|An FC port is the physical port of an FC adapter on a cluster node that  
can be connected to an FC network.
```

```
|is_home
```

```
|boolean
```

```
a|Indicates whether or not the FC interface currently resides on the home  
node.
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|port
```

```
|link:#fc_port_reference[fc_port_reference]
```

```
a|An FC port is the physical port of an FC adapter on a cluster node that  
can be connected to an FC network.
```

```
|===
```

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.


```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

latency

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#throughput]

[.api-collapsible-fifth-title]

throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

Performance numbers, such as IOPS latency and throughput

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops
|link:#iops[iops]

a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status
|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with

"backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#iops_raw]

[.api-collapsible-fifth-title]

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]  
[.api-collapsible-fifth-title]  
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description  
  
|read  
|integer  
a|Performance metric for read I/O operations.
```

```
|total  
|integer  
a|Performance metric aggregated over all types of I/O operations.
```

```
|write  
|integer  
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]  
[.api-collapsible-fifth-title]  
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.

|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]  
[.api-collapsible-fifth-title]  
svm
```

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|name  
|string  
a|The name of the SVM.
```

```
|uuid  
|string  
a|The unique identifier of the SVM.
```

```
|===
```

```
[#fc_interface]  
[.api-collapsible-fifth-title]  
fc_interface
```

A Fibre Channel (FC) interface is the logical endpoint for FC network connections to an SVM. An FC interface provides FC access to storage within the interface SVM using either Fibre Channel Protocol or NVMe over Fibre Channel (NVMe/FC).

An FC interface is created on an FC port which is located on a cluster node. The FC port must be specified to identify the location of the interface for a POST or PATCH operation that relocates an interface. You can identify the port by supplying either the node and port names or the port UUID.

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|comment
|string
a|A user configurable comment. Optional in POST; valid in PATCH. To clear
a prior comment, set the property to an empty string in PATCH.

|data_protocol
|string
a|The data protocol for which the FC interface is configured. Required in
POST.

|enabled
|boolean
a|The administrative state of the FC interface. The FC interface can be
disabled to block all FC communication with the SVM through this
interface. Optional in POST and PATCH; defaults to _true_ (enabled) in
POST.

|location
|link:#location[location]
a|The location of the FC interface is defined by the location of its port.
An FC port is identified by its UUID, or a combination of its node name
and port name. Either the UUID or the node name and port name are required
for POST. To move an interface, supply either the UUID or the node name
and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and
"location.home_port" during a POST or PATCH. "location.node" and
"location.port" refer to the current location of the FC interface. This
may be different from "location.home_node" and "location.home_port" in
instances where the FC interface failed over due to an offline node.

|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput

```

|name
|string
a|The name of the FC interface. Required in POST; optional in PATCH.

|port_address
|string
a|The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a hexadecimal encoded numeric value.

|state
|string
a|The current operational state of the FC interface. The state is set to `_down_` if the interface is not enabled.

If the node hosting the port is down or unavailable, no state value is returned.

|statistics
|link:#statistics[statistics]
a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|svm
|link:#svm[svm]
a|

|uuid
|string
a|The unique identifier of the FC interface. Required in the URL.

|wwnn

```
|string
a|The world wide node name (WWNN) of the FC interface SVM. The WWNN is
generated by ONTAP when Fibre Channel Protocol or the NVMe service is
created for the FC interface SVM.
```

```
* example: 20:00:00:50:56:b4:13:01
* readOnly: 1
* Introduced in: 9.6
```

```
|wwpn
|string
a|The world wide port name (WWPN) of the FC interface. The WWPN is
generated by ONTAP when the FC interface is created.
```

```
* example: 20:00:00:50:56:b4:13:a8
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
```

```

[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDa5001084b051d034fb6d192dc5ce0053]]
= Create an FC interface

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/network/fc/interfaces`#

*Introduced In:* 9.6

Creates an FC interface.

```

== Required properties

- * ``svm.uuid`` or ``svm.name`` - Existing SVM in which to create the FC interface.
- * ``name`` - Name of the FC interface.
- * ``location.port.uuid`` or both ``location.port.name`` and ``location.port.node.name`` - FC port on which to create the FC interface.
- * ``data_protocol`` - Data protocol for the FC interface.

== Default property values

If not specified in POST, the following default property values are assigned.

- * ``enabled`` - `_true_`

== Related ONTAP commands

- * ``network interface create``

== Learn more

- * `xref:{relative_path}network_fc_interfaces_endpoint_overview.html[DOC/network/fc/interfaces]`

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

- * Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|comment

|string

a|A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH.

|data_protocol

|string

a|The data protocol for which the FC interface is configured. Required in POST.

|enabled

|boolean

a|The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to `_true_` (enabled) in POST.

|location

|link:#location[location]

a|The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its node name and port name. Either the UUID or the node name and port name are required for POST. To move an interface, supply either the UUID or the node name and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

|name

|string

a|The name of the FC interface. Required in POST; optional in PATCH.

|port_address

|string

a|The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a hexadecimal encoded numeric value.

|state

|string

a|The current operational state of the FC interface. The state is set to `_down_` if the interface is not enabled.

If the node hosting the port is down or unavailable, no state value is returned.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The unique identifier of the FC interface. Required in the URL.

|wwnn

|string

a|The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.

* example: 20:00:00:50:56:b4:13:01

* readOnly: 1

* Introduced in: 9.6

|wwpn

```
|string
a|The world wide port name (WWPN) of the FC interface. The WWPN is
generated by ONTAP when the FC interface is created.

* example: 20:00:00:50:56:b4:13:a8
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "data_protocol": "string",
  "location": {
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "lif1",
  "port_address": "5060F",
  "state": "string",
```



```

"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"wwnn": "20:00:00:50:56:b4:13:01",
"wwpn": "20:00:00:50:56:b4:13:a8"
}
=====

```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#fc_interface[fc_interface]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "comment": "string",
      "data_protocol": "string",
      "location": {
        "home_node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    },

```

```

    "home_port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "lif1",
  "port_address": "5060F",
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "wwnn": "20:00:00:50:56:b4:13:01",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 1966140
| An interface with the same name already exists.

```

```
| 1966217
| The specified port is not valid on the node provided.

| 2621462
| The supplied SVM does not exist.

| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5373966
| A Fibre Channel interface with the _fcp_ protocol cannot be created in
an SVM that is configured for NVMe.

| 5374102
| The specified Fibre Channel interface cannot be created because the
Fibre Channel adapter is down. Bring the adapter up and try again.

| 5374871
| The Fibre Channel port identified by the specified UUID does not refer
to the same port as that identified by the specified node name and/or port
name.

| 5374872
| If either `location.port.node.name` or `location.port.name` is supplied,
both properties must be supplied.

| 5374873
| The Fibre Channel port must be specified using either
`location.port.uuid` or `location.port.node.name` and
`location.port.name`.

| 72089652
| An NVMe service must be created before creating a Fibre Channel
interface using the NVMe over FC data protocol.

| 72089672
| The specified Fibre Channel port does not support the NVMe over FC data
protocol.

| 72089900
| A Fibre Channel interface with the _fc_nvme_ protocol cannot be created
in an SVM that is configured for a SAN protocol.
|===
```

```

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#home_node]
[.api-collapsible-fifth-title]
home_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

```

```
[#node]
[.api-collapsible-fifth-title]
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#fc_port_reference]
[.api-collapsible-fifth-title]
fc_port_reference
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the FC port.
```

```
|node
|link:#node[node]
a|The node on which the FC port is located.
```

```
|uuid
|string
a|The unique identifier of the FC port.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its node name and port name. Either the UUID or the node name and port name are required for POST. To move an interface, supply either the UUID or the node name and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|home_node
|link:#home_node[home_node]
a|

|home_port
|link:#fc_port_reference[fc_port_reference]
a|An FC port is the physical port of an FC adapter on a cluster node that
can be connected to an FC network.

|is_home
|boolean
a|Indicates whether or not the FC interface currently resides on the home
node.

|node
|link:#node[node]
a|

|port
|link:#fc_port_reference[fc_port_reference]
a|An FC port is the physical port of an FC adapter on a cluster node that
can be connected to an FC network.

|===

[#iops]
[.api-collapsible-fifth-title]
iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```



```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
```

metric

Performance numbers, such as IOPS latency and throughput

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|duration
```

```
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
```

```
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
```

```
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
```

```
|link:#throughput[throughput]
```

a|The rate of throughput bytes per second observed at the storage object.

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
```

```
|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#fc_interface]
[.api-collapsible-fifth-title]
fc_interface
```

A Fibre Channel (FC) interface is the logical endpoint for FC network connections to an SVM. An FC interface provides FC access to storage within the interface SVM using either Fibre Channel Protocol or NVMe over Fibre Channel (NVMe/FC).

An FC interface is created on an FC port which is located on a cluster node. The FC port must be specified to identify the location of the interface for a POST or PATCH operation that relocates an interface. You can identify the port by supplying either the node and port names or the port UUID.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|comment
|string
```

a|A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH.

```
|data_protocol
|string
```

a|The data protocol for which the FC interface is configured. Required in POST.

|enabled

|boolean

a|The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to `_true_` (enabled) in POST.

|location

|link:#location[location]

a|The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its node name and port name. Either the UUID or the node name and port name are required for POST. To move an interface, supply either the UUID or the node name and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

|name

|string

a|The name of the FC interface. Required in POST; optional in PATCH.

|port_address

|string

a|The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a hexadecimal encoded numeric value.

|state

|string

a|The current operational state of the FC interface. The state is set to `_down_` if the interface is not enabled.

If the node hosting the port is down or unavailable, no state value is returned.

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|The unique identifier of the FC interface. Required in the URL.
```

```
|wwnn
|string
a|The world wide node name (WWNN) of the FC interface SVM. The WWNN is
generated by ONTAP when Fibre Channel Protocol or the NVMe service is
created for the FC interface SVM.
```

```
* example: 20:00:00:50:56:b4:13:01
* readOnly: 1
* Introduced in: 9.6
```

```
|wwpn
|string
a|The world wide port name (WWPN) of the FC interface. The WWPN is
generated by ONTAP when the FC interface is created.
```

```
* example: 20:00:00:50:56:b4:13:a8
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====

[[ID25b9eeec11f3b35897408b1d4ac78423]]
= Delete an FC interface

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/network/fc/interfaces/{uuid}`#

*Introduced In:* 9.6

Deletes an FC interface.

== Related ONTAP commands

* `network interface delete`

== Learn more

* xref:{relative_path}network_fc_interfaces_endpoint_overview.html[DOC
/network/fc/interfaces]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|The unique identifier for the FC interface.

|===

== Response
```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

===

Error Code	Description
53280992	The FC interface could not be deleted because it is enabled.

===

[cols=3*,options=header]

===

Name	Type	Description
error		
link:#error[error]		

a|

===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

{
 "error": {
 "arguments": [
 {
 "code": "string",
 "message": "string"
 }
],
 "code": "4",
 "message": "entry doesn't exist",
 "target": "uuid"
 }
}

```

    }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

```

```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```

|===

```

```

//end collapsible .Definitions block
====

```

```

[[ID0ffcead2844e2ab9c4c1676647697d10]]
= Retrieve an FC interface

```

```

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/fc/interfaces/{uuid}`#

```

Introduced In: 9.6

Retrieves an FC interface.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```

* `statistics.+++`
* `metric.+++`

```

== Related ONTAP commands

```
* `network interface show`  
* `vserver fcp interface show`
```

== Learn more

```
* xref:{relative_path}network_fc_interfaces_endpoint_overview.html[DOC  
/network/fc/interfaces]
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|uuid  
|string  
|path  
|True  
a|The unique identifier for the FC interface.
```

```
|fields  
|array[string]  
|query  
|False  
a|Specify the fields to return.
```

```
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|comment
|string
a|A user configurable comment. Optional in POST; valid in PATCH. To clear
a prior comment, set the property to an empty string in PATCH.
```

```
|data_protocol
|string
a|The data protocol for which the FC interface is configured. Required in
POST.
```

```
|enabled
|boolean
a|The administrative state of the FC interface. The FC interface can be
disabled to block all FC communication with the SVM through this
interface. Optional in POST and PATCH; defaults to _true_ (enabled) in
POST.
```

```
|location
|link:#location[location]
a|The location of the FC interface is defined by the location of its port.
An FC port is identified by its UUID, or a combination of its node name
and port name. Either the UUID or the node name and port name are required
for POST. To move an interface, supply either the UUID or the node name
and port name in a PATCH.
```

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

```
|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput
```

```
|name
|string
a|The name of the FC interface. Required in POST; optional in PATCH.
```


|port_address

|string

a|The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a hexadecimal encoded numeric value.

|state

|string

a|The current operational state of the FC interface. The state is set to `_down_` if the interface is not enabled.

If the node hosting the port is down or unavailable, no state value is returned.

|statistics

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The unique identifier of the FC interface. Required in the URL.

|wwnn

|string

a|The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.

```
* example: 20:00:00:50:56:b4:13:01
```

```
* readOnly: 1
```

```
* Introduced in: 9.6
```

```
|wwpn
```

```
|string
```

```
a|The world wide port name (WWPN) of the FC interface. The WWPN is  
generated by ONTAP when the FC interface is created.
```

```
* example: 20:00:00:50:56:b4:13:a8
```

```
* readOnly: 1
```

```
* Introduced in: 9.6
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{  
  "_links": {  
    "self": {  
      "href": "/api/resourcelink"  
    }  
  },  
  "comment": "string",  
  "data_protocol": "string",  
  "location": {  
    "home_node": {  
      "_links": {  
        "self": {  
          "href": "/api/resourcelink"  
        }  
      },  
      "name": "node1",  
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"  
    },  
    "home_port": {  
      "_links": {  
        "self": {  
          "href": "/api/resourcelink"  
        }  
      },  
    },  
  },  
}
```

```

    "name": "0a",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "port": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "0a",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",

```

```

    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "name": "lif1",
  "port_address": "5060F",
  "state": "string",
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "wwnn": "20:00:00:50:56:b4:13:01",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
====

== Error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#home_node]
[.api-collapsible-fifth-title]
home_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

```

```
|uuid
|string
a|
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#fc_port_reference]
[.api-collapsible-fifth-title]
fc_port_reference
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
```

```
|string
a|The name of the FC port.
```

```
|node
|link:#node[node]
a|The node on which the FC port is located.
```

```
|uuid
|string
a|The unique identifier of the FC port.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#location]
```

```
[.api-collapsible-fifth-title]
```

```
location
```

The location of the FC interface is defined by the location of its port.

An FC port is identified by its UUID, or a combination of its node name and port name. Either the UUID or the node name and port name are required for POST. To move an interface, supply either the UUID or the node name and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|home_node
```

```
|link:#home_node[home_node]
```

```
a|
```

```
|home_port
```

```
|link:#fc_port_reference[fc_port_reference]
```

```
a|An FC port is the physical port of an FC adapter on a cluster node that  
can be connected to an FC network.
```

```
|is_home
```

```
|boolean
```

```
a|Indicates whether or not the FC interface currently resides on the home  
node.
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|port
```

```
|link:#fc_port_reference[fc_port_reference]
```

```
a|An FC port is the physical port of an FC adapter on a cluster node that  
can be connected to an FC network.
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
```

```

|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer

```

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#metric]

[.api-collapsible-fifth-title]

metric

Performance numbers, such as IOPS latency and throughput

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|duration

|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops

|link:#iops[iops]

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#iops_raw]

[.api-collapsible-fifth-title]

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
```

```
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.


```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
```

```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

[[ID7ff11a4e38ab39360f77e3ecbd16317e]]

= Update an FC interface

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-block]#`/network/fc/interfaces/{uuid}`#

Introduced In: 9.6

Updates an FC interface.

== Related ONTAP commands

* `network interface modify`

== Learn more

* xref:{relative_path}network_fc_interfaces_endpoint_overview.html[DOC /network/fc/interfaces]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|The unique identifier for the FC interface.

|===

== Request Body

[cols=3*,options=header]

|===

|Name
|Type
|Description

|comment
|string

a|A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH.

|enabled
|boolean

a|The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to `_true_` (enabled) in POST.

|location
|link:#location[location]

a|The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its node name and port name. Either the UUID or the node name and port name are required for POST. To move an interface, supply either the UUID or the node name and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

|name
|string

a|The name of the FC interface. Required in POST; optional in PATCH.

|port_address
|string

a|The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).

This is useful for obtaining statistics and diagnostic information from FC

switches.

This is a hexadecimal encoded numeric value.

|state

|string

a|The current operational state of the FC interface. The state is set to `_down_` if the interface is not enabled.

If the node hosting the port is down or unavailable, no state value is returned.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The unique identifier of the FC interface. Required in the URL.

|wwnn

|string

a|The world wide node name (WWNN) of the FC interface SVM. The WWNN is generated by ONTAP when Fibre Channel Protocol or the NVMe service is created for the FC interface SVM.

* example: 20:00:00:50:56:b4:13:01

* readOnly: 1

* Introduced in: 9.6

|wwpn

|string

a|The world wide port name (WWPN) of the FC interface. The WWPN is generated by ONTAP when the FC interface is created.

* example: 20:00:00:50:56:b4:13:a8

* readOnly: 1

* Introduced in: 9.6

|===

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "comment": "string",
  "location": {
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "name": "0a",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "lif1",
  "port_address": "5060F",
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "wwnn": "20:00:00:50:56:b4:13:01",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
====

== Response
```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

1966140	An interface with the same name already exists.
---------	-------------------------------------------------

1966217	The specified port is not valid on the node provided.
---------	-------------------------------------------------------

1966238	The node or port of an active SAN data interface cannot be changed.
---------	---------------------------------------------------------------------

1966702	The destination node is not healthy.
---------	--------------------------------------

5374579	The SAN Kernel Agent on the node is unavailable.
---------	--------------------------------------------------

5374870	A partial failure occurred; renaming the interface failed. Correct the error and resubmit the request.
---------	--------------------------------------------------------------------------------------------------------

5374871	The Fibre Channel port identified by the specified UUID does not refer to the same port as that identified by the specified node name and/or port name.
---------	---------------------------------------------------------------------------------------------------------------------------------------------------------

5374872	If either `location.port.node.name` or `location.port.name` is supplied, both properties must be supplied.
---------	------------------------------------------------------------------------------------------------------------

72089674	You cannot move a Fibre Channel interface configured for the NVMe over FC data protocol.
----------	------------------------------------------------------------------------------------------

|===

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#home_node]
[.api-collapsible-fifth-title]
home_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#node]

```



```
[.api-collapsible-fifth-title]
```

```
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#fc_port_reference]
```

```
[.api-collapsible-fifth-title]
```

```
fc_port_reference
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the FC port.
```

```
|node
```

```
|link:#node[node]
```

```
a|The node on which the FC port is located.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the FC port.
```

```
|===
```

```
[#node]  
[.api-collapsible-fifth-title]  
node
```

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|name  
|string  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#location]  
[.api-collapsible-fifth-title]  
location
```

The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its node name and port name. Either the UUID or the node name and port name are required for POST. To move an interface, supply either the UUID or the node name and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

```
[cols=3*,options=header]  
|===  
|Name  
|Type
```

|Description

|home_node

|link:#home_node[home_node]

a|

|home_port

|link:#fc_port_reference[fc_port_reference]

a|An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

|is_home

|boolean

a|Indicates whether or not the FC interface currently resides on the home node.

|node

|link:#node[node]

a|

|port

|link:#fc_port_reference[fc_port_reference]

a|An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

|===

[#iops]

[.api-collapsible-fifth-title]

iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be

metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
```

```
[.api-collapsible-fifth-title]
```

```
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

```
metric
```

Performance numbers, such as IOPS latency and throughput

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string

```

a|The timestamp of the performance data.

|===

[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read
|integer

a|Performance metric for read I/O operations.

|total
|integer

a|Performance metric aggregated over all types of I/O operations.

|write
|integer

a|Performance metric for write I/O operations.

|===

[#latency_raw]
[.api-collapsible-fifth-title]

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

```

```

[#statistics]
[.api-collapsible-fifth-title]
statistics

```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw
|link:#latency_raw[latency_raw]

```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#fc_interface]

[.api-collapsible-fifth-title]

fc_interface

A Fibre Channel (FC) interface is the logical endpoint for FC network connections to an SVM. An FC interface provides FC access to storage within the interface SVM using either Fibre Channel Protocol or NVMe over Fibre Channel (NVMe/FC).

An FC interface is created on an FC port which is located on a cluster node. The FC port must be specified to identify the location of the interface for a POST or PATCH operation that relocates an interface. You can identify the port by supplying either the node and port names or the port UUID.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|comment

|string

a|A user configurable comment. Optional in POST; valid in PATCH. To clear a prior comment, set the property to an empty string in PATCH.

|enabled

|boolean

a|The administrative state of the FC interface. The FC interface can be disabled to block all FC communication with the SVM through this interface. Optional in POST and PATCH; defaults to `_true_` (enabled) in POST.

|location

|link:#location[location]

a|The location of the FC interface is defined by the location of its port. An FC port is identified by its UUID, or a combination of its node name and port name. Either the UUID or the node name and port name are required for POST. To move an interface, supply either the UUID or the node name and port name in a PATCH.

The location of an FC interface can be set using "location.home_node" and "location.home_port" during a POST or PATCH. "location.node" and "location.port" refer to the current location of the FC interface. This may be different from "location.home_node" and "location.home_port" in instances where the FC interface failed over due to an offline node.

|name

|string

a|The name of the FC interface. Required in POST; optional in PATCH.

|port_address

|string

a|The port address of the FC interface. Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the FC interface when the SVM performs a port login (PLOGI).

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a hexadecimal encoded numeric value.

|state

|string

a|The current operational state of the FC interface. The state is set to `_down_` if the interface is not enabled.

If the node hosting the port is down or unavailable, no state value is returned.

|svm

|link:#svm[svm]

a|

```

|uuid
|string
a|The unique identifier of the FC interface. Required in the URL.

|wwnn
|string
a|The world wide node name (WWNN) of the FC interface SVM. The WWNN is
generated by ONTAP when Fibre Channel Protocol or the NVMe service is
created for the FC interface SVM.

* example: 20:00:00:50:56:b4:13:01
* readOnly: 1
* Introduced in: 9.6

|wwpn
|string
a|The world wide port name (WWPN) of the FC interface. The WWPN is
generated by ONTAP when the FC interface is created.

* example: 20:00:00:50:56:b4:13:a8
* readOnly: 1
* Introduced in: 9.6

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID431d9c277f2171853af6f8c2c0d7dea5]]
= Retrieve FC interface historical performance metrics

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
```

```
block]#`/network/fc/interfaces/{uuid}/metrics`#
```

Introduced In: 9.8

Retrieves historical performance metrics for an FC interface.

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|throughput.write
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by throughput.write
```

```
|throughput.read
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by throughput.read
```

```
|throughput.total
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by throughput.total
```

```
|duration
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by duration
```

```
|timestamp
```

```
|string  
|query  
|False  
a|Filter by timestamp
```

```
|status  
|string  
|query  
|False  
a|Filter by status
```

```
|iops.total  
|integer  
|query  
|False  
a|Filter by iops.total
```

```
|iops.read  
|integer  
|query  
|False  
a|Filter by iops.read
```

```
|iops.other  
|integer  
|query  
|False  
a|Filter by iops.other
```

```
|iops.write  
|integer  
|query  
|False  
a|Filter by iops.write
```

```
|latency.total  
|integer  
|query  
|False  
a|Filter by latency.total
```



```
|latency.read
|integer
|query
|False
a|Filter by latency.read
```

```
|latency.other
|integer
|query
|False
a|Filter by latency.other
```

```
|latency.write
|integer
|query
|False
a|Filter by latency.write
```

```
|uuid
|string
|path
|True
a|Unique identifier of the FC interface.
```

```
|interval
|string
|query
|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:
```

- * 1h: Metrics over the most recent hour sampled over 15 seconds.
- * 1d: Metrics over the most recent day sampled over 5 minutes.
- * 1w: Metrics over the most recent week sampled over 30 minutes.
- * 1m: Metrics over the most recent month sampled over 2 hours.
- * 1y: Metrics over the most recent year sampled over a day.
- * Default value: 1
- * enum: ["1h", "1d", "1w", "1m", "1y"]

```
|return_timeout
|integer
```

```

|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|===

== Response

```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#records[records]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
```

```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
    }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [

```

```

        {
            "code": "string",
            "message": "string"
        }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
}
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

|===

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|next
|link:href[href]
a|
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#records]
[.api-collapsible-fifth-title]
records
```

Performance numbers, such as IOPS latency and throughput

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp

```

```
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
```

```

a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```

|===

```

```

//end collapsible .Definitions block
====

```

```

:leveloffset: -1

```

```

= Retrieve FC port information

```

```

:leveloffset: +1

```

```

[[ID7e368a1567d2503fad49a85cdd39a3f9]]
= Network FC ports endpoint overview

```

```

== Overview

```

Fibre Channel (FC) ports are the physical ports of FC adapters on ONTAP cluster nodes that can be connected to FC networks to provide FC network connectivity. An FC port defines the location of an FC interface within the ONTAP cluster.

The Fibre Channel port REST API allows you to discover FC ports, obtain status information for FC ports, and configure FC port properties. POST and DELETE requests are not supported. You must physically add and remove FC adapters to ONTAP nodes to create and remove ports from the ONTAP cluster.

== Performance monitoring

Performance of an FC port can be monitored by observing the ``metric.++`` and ``statistics.++`` properties. These properties show the performance of an FC port in terms of IOPS, latency, and throughput. The ``metric.++`` properties denote an average, whereas ``statistics.++`` properties denote a real-time monotonically increasing value aggregated across all nodes.

== Examples

=== Retrieving all FC ports

The API:

GET /api/network/fc/ports

The call:

```
curl -X GET "https://<mgmt-ip>/api/network/fc/ports" -H "accept:
application/hal+json"
```

The response:

```
{
  "records": [
    {
      "node": {
        "name": "node1",
        "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
          }
        }
      },
      "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
      "name": "0a",
      "_links": {
        "self": {
          "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-
005056bb838e"
        }
      }
    },
    {
      "node": {
```

```

    "name": "node1",
    "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-629ceb62a497"
      }
    }
  },
  "uuid": "931b23f7-b047-11e8-9af3-005056bb838e",
  "name": "0b",
  "_links": {
    "self": {
      "href": "/api/network/fc/ports/931b23f7-b047-11e8-9af3-005056bb838e"
    }
  }
},
{
  "node": {
    "name": "node1",
    "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-629ceb62a497"
      }
    }
  },
  "uuid": "931b25ba-b047-11e8-9af3-005056bb838e",
  "name": "0c",
  "_links": {
    "self": {
      "href": "/api/network/fc/ports/931b25ba-b047-11e8-9af3-005056bb838e"
    }
  }
},
{
  "node": {
    "name": "node1",
    "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-629ceb62a497"
      }
    }
  }
}

```

```

    }
  }
},
"uuid": "931b2748-b047-11e8-9af3-005056bb838e",
"name": "0d",
"_links": {
  "self": {
    "href": "/api/network/fc/ports/931b2748-b047-11e8-9af3-005056bb838e"
  }
}
},
{
  "node": {
    "name": "node1",
    "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-629ceb62a497"
      }
    }
  },
  "uuid": "931b28c2-b047-11e8-9af3-005056bb838e",
  "name": "0e",
  "_links": {
    "self": {
      "href": "/api/network/fc/ports/931b28c2-b047-11e8-9af3-005056bb838e"
    }
  }
},
{
  "node": {
    "name": "node1",
    "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-629ceb62a497"
      }
    }
  },
  "uuid": "931b2a7b-b047-11e8-9af3-005056bb838e",
  "name": "0f",
  "_links": {

```

```

        "self": {
            "href": "/api/network/fc/ports/931b2a7b-b047-11e8-9af3-005056bb838e"
        }
    },
    {
        "node": {
            "name": "node1",
            "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
            "_links": {
                "self": {
                    "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-629ceb62a497"
                }
            }
        },
        "uuid": "931b2e2b-b047-11e8-9af3-005056bb838e",
        "name": "1b",
        "_links": {
            "self": {
                "href": "/api/network/fc/ports/931b2e2b-b047-11e8-9af3-005056bb838e"
            }
        }
    }
],
"num_records": 8,
"_links": {
    "self": {
        "href": "/api/network/fc/ports"
    }
}
}

```

'''

=== Retrieving all FC ports with state `_online_`

The ``state`` query parameter is used to perform the query.

The API:

GET /api/network/fc/ports

```
# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/ports?state=online" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "node": {
        "name": "node1",
        "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
          }
        }
      },
      "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
      "name": "0a",
      "state": "online",
      "_links": {
        "self": {
          "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-
005056bb838e"
        }
      }
    },
    {
      "node": {
        "name": "node1",
        "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-
629ceb62a497"
          }
        }
      },
      "uuid": "931b23f7-b047-11e8-9af3-005056bb838e",
      "name": "0b",
      "state": "online",
      "_links": {
        "self": {
          "href": "/api/network/fc/ports/931b23f7-b047-11e8-9af3-
005056bb838e"
        }
      }
    }
  ]
}
```



```

    }
  },
  {
    "node": {
      "name": "node1",
      "uuid": "3c768e01-1abc-4b3b-b7c0-629ceb62a497",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3c768e01-1abc-4b3b-b7c0-629ceb62a497"
        }
      }
    },
    "uuid": "931b25ba-b047-11e8-9af3-005056bb838e",
    "name": "0c",
    "state": "online",
    "_links": {
      "self": {
        "href": "/api/network/fc/ports/931b25ba-b047-11e8-9af3-005056bb838e"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/network/fc/ports?state=online"
  }
}
}

```

'''

=== Retrieving an FC port

The API:

GET /api/network/fc/ports/{uuid}

The call:

curl -X GET "https://<mgmt-ip>/api/network/fc/ports/931b20f8-b047-11e8-9af3-005056bb838e" -H "accept: application/hal+json"

```
# The response:
{
  "node": {
    "name": "node1",
    "uuid": "5a534a72-b047-11e8-9af3-005056bb838e",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/5a534a72-b047-11e8-9af3-005056bb838e"
      }
    }
  },
  "uuid": "931b20f8-b047-11e8-9af3-005056bb838e",
  "name": "0a",
  "description": "Fibre Channel Target Adapter 0a (ACME Fibre Channel Adapter, rev. 1.0.0, 8G)",
  "enabled": true,
  "fabric": {
    "connected": true,
    "connected_speed": 8,
    "name": "55:0e:b1:a0:20:40:80:00",
    "port_address": "52100",
    "switch_port": "ssan-g620-03:1"
  },
  "physical_protocol": "fibre_channel",
  "speed": {
    "maximum": "8",
    "configured": "auto"
  },
  "state": "online",
  "supported_protocols": [
    "fc"
  ],
  "transceiver": {
    "form_factor": "SFP",
    "manufacturer": "ACME",
    "capabilities": [
      4,
      8
    ],
    "part_number": "1000"
  },
  "wwnn": "50:0a:09:80:bb:83:8e:00",
  "wwpn": "50:0a:09:82:bb:83:8e:00",
  "metric": {
    "timestamp": "2019-04-09T05:50:15Z",
    "duration": "PT15S",
  }
}
```

```
"status": "ok",
"latency": {
  "other": 0,
  "total": 0,
  "read": 0,
  "write": 0
},
"iops": {
  "read": 0,
  "write": 0,
  "other": 0,
  "total": 0
},
"throughput": {
  "read": 0,
  "write": 0,
  "total": 0
}
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "total": 0
  }
},
"_links": {
  "self": {
    "href": "/api/network/fc/ports/931b20f8-b047-11e8-9af3-005056bb838e"
  }
}
}
```

'''

=== Disabling an FC port

If an active FC interface exists on an FC port, the port cannot be disabled.

The API:

PATCH /api/network/fc/ports/{uuid}

The call:

```
curl -X PATCH "http://<mgmt-ip>/api/network/fc/ports/931b20f8-b047-11e8-9af3-005056bb838e" -H "accept: application/hal+json" -d '{ "enabled": false }'
```

[[IDffffde6ac93b8abbedd55a23fa377dd9f]]

= Retrieve FC ports

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/network/fc/ports`#

Introduced In: 9.6

Retrieves FC ports.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `fabric.name`

* `statistics.+++`

* `metric.+++`

== Related ONTAP commands

* `network fcp adapter show`

== Learn more

* xref:{relative_path}network_fc_ports_endpoint_overview.html[DOC
/network/fc/ports]

== Parameters

[cols=5*,options=header]
|==

|Name
|Type
|In
|Required
|Description

|description
|string
|query
|False
a|Filter by description

|name
|string
|query
|False
a|Filter by name

|metric.latency.total
|integer
|query
|False
a|Filter by metric.latency.total

* Introduced in: 9.8

|metric.latency.read
|integer
|query
|False

```
a|Filter by metric.latency.read
```

```
* Introduced in: 9.8
```

```
|metric.latency.other
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by metric.latency.other
```

```
* Introduced in: 9.8
```

```
|metric.latency.write
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by metric.latency.write
```

```
* Introduced in: 9.8
```

```
|metric.status
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by metric.status
```

```
* Introduced in: 9.8
```

```
|metric.iops.total
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by metric.iops.total
```

```
* Introduced in: 9.8
```

```
|metric.iops.read
```

```
|integer
```

```
|query
```

```
|False
```

```
a|Filter by metric.iops.read
```

* Introduced in: 9.8

```
|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other
```

* Introduced in: 9.8

```
|metric.iops.write
|integer
|query
|False
a|Filter by metric.iops.write
```

* Introduced in: 9.8

```
|metric.duration
|string
|query
|False
a|Filter by metric.duration
```

* Introduced in: 9.8

```
|metric.timestamp
|string
|query
|False
a|Filter by metric.timestamp
```

* Introduced in: 9.8

```
|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write
```

* Introduced in: 9.8

```
|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read
```

* Introduced in: 9.8

```
|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total
```

* Introduced in: 9.8

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|wwnn
|string
|query
|False
a|Filter by wwnn
```

```
|wwpn
|string
|query
|False
a|Filter by wwpn
```

```
|node.name
|string
|query
|False
a|Filter by node.name
```



```
|node.uuid  
|string  
|query  
|False  
a|Filter by node.uuid
```

```
|physical_protocol  
|string  
|query  
|False  
a|Filter by physical_protocol
```

```
|fabric.name  
|string  
|query  
|False  
a|Filter by fabric.name
```

```
|fabric.port_address  
|string  
|query  
|False  
a|Filter by fabric.port_address
```

```
|fabric.connected  
|boolean  
|query  
|False  
a|Filter by fabric.connected
```

```
|fabric.connected_speed  
|integer  
|query  
|False  
a|Filter by fabric.connected_speed
```

```
|fabric.switch_port  
|string  
|query  
|False  
a|Filter by fabric.switch_port
```

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|speed.maximum
|string
|query
|False
a|Filter by speed.maximum
```

```
|speed.configured
|string
|query
|False
a|Filter by speed.configured
```

```
|supported_protocols
|string
|query
|False
a|Filter by supported_protocols
```

```
|state
|string
|query
|False
a|Filter by state
```

```
|statistics.iops_raw.total
|integer
|query
|False
a|Filter by statistics.iops_raw.total
```

* Introduced in: 9.8

```
|statistics.iops_raw.read
```

```
|integer
|query
|False
a|Filter by statistics.iops_raw.read
```

* Introduced in: 9.8

```
|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other
```

* Introduced in: 9.8

```
|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write
```

* Introduced in: 9.8

```
|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total
```

* Introduced in: 9.8

```
|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read
```

* Introduced in: 9.8

```
|statistics.latency_raw.other
|integer
|query
```

```

|False
a|Filter by statistics.latency_raw.other

* Introduced in: 9.8

|statistics.latency_raw.write
|integer
|query
|False
a|Filter by statistics.latency_raw.write

* Introduced in: 9.8

|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp

* Introduced in: 9.8

|statistics.status
|string
|query
|False
a|Filter by statistics.status

* Introduced in: 9.8

|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write

* Introduced in: 9.8

|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read

```

* Introduced in: 9.8

```
|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total
```

* Introduced in: 9.8

```
|transceiver.manufacturer
|string
|query
|False
a|Filter by transceiver.manufacturer
```

```
|transceiver.part_number
|string
|query
|False
a|Filter by transceiver.part_number
```

```
|transceiver.form-factor
|string
|query
|False
a|Filter by transceiver.form-factor
```

```
|transceiver.capabilities
|integer
|query
|False
a|Filter by transceiver.capabilities
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

```

```

|records
|array[link:#fc_port[fc_port]]
a|

```

```

|===

```

```

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "description": "Fibre Channel Target Adapter 0a (ACME Fibre Channel
Adapter, rev. 1.0.0, 8G)",
      "fabric": {
        "connected_speed": "16",
        "name": "string",

```

```

    "port_address": "52100A",
    "switch_port": "ssan-g620-03:33"
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "name": "0a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "physical_protocol": "string",
  "speed": {
    "configured": "auto",
    "maximum": "32"
  },
  "state": "online",
  "statistics": {
    "iops_raw": {

```



```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"supported_protocols": [
    "string"
],
"transceiver": {
    "capabilities": [
        "16"
    ],
    "form-factor": "string",
    "manufacturer": "Acme, Inc.",
    "part_number": "string"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"wwnn": "20:00:00:50:56:b4:13:a8",
"wwpn": "20:00:00:50:56:b4:13:a8"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string

```

```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#fabric]
[.api-collapsible-fifth-title]
fabric

```

Properties of the fabric to which the FC port is attached.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|connected
```

```
|boolean
```

a|Reports if the physical port has established a connection with the FC fabric.

```
|connected_speed
```

```
|integer
```

a|The negotiated data rate between the target FC port and the fabric in gigabits per second.

```
|name
```

```
|string
```

a|The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|port_address
```

```
|string
```

a|The FC port address of the host bus adapter (HBA) physical port.

Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the FC port address given to the physical host bus adapter (HBA) port when the port performs a fabric login (FLOGI).

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a six-digit hexadecimal encoded numeric value.

```
|switch_port
|string
a|The switch port to which the FC port is connected.
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS latency and throughput

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#node]

[.api-collapsible-fifth-title]

node

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links


```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#speed]
```

```
[.api-collapsible-fifth-title]
```

```
speed
```

The physical device speed related properties of the FC port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|configured
```

```
|string
```

```
a|The configured speed of the FC port in gigabits per second.
```

```
|maximum
```

```
|string
```

```
a|The maximum speed supported by the FC port in gigabits per second.
```

```
|===
```

```
[#iops_raw]
```

```
[.api-collapsible-fifth-title]
```

```
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be  
metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
```

```
[.api-collapsible-fifth-title]
```

```
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
```

```
|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time
```

between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#transceiver]
[.api-collapsible-fifth-title]
transceiver
```

Properties of the transceiver connected to the FC port.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|capabilities
|array[integer]
a|The speeds of which the transceiver is capable in gigabits per second.
```

```
|form-factor
|string
a|The form factor of the transceiver. Possible values are:
```

```
* _sfp_ - Small Form Factor - Pluggable
* _sff_ - Small Form Factor
* _unk_ - Unknown
```

```
|manufacturer
|string
a|The manufacturer of the transceiver.
```

```
|part_number
|string
a|The part number of the transceiver.
```

```
|===
```

```
[#fc_port]
[.api-collapsible-fifth-title]
fc_port
```

A Fibre Channel (FC) port is the physical port of an FC adapter on an ONTAP cluster node that can be connected to an FC network to provide FC network connectivity. An FC port defines the location of an FC interface within the ONTAP cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|description
|string
a|A description of the FC port.
```

```
|enabled
|boolean
a|The administrative state of the FC port. If this property is set to _false_, all FC connectivity to FC interfaces are blocked. Optional in PATCH.
```

```
|fabric
|link:#fabric[fabric]
a|Properties of the fabric to which the FC port is attached.
```

```
|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput
```

```
|name
|string
a|The FC port name.
```

```
|node
|link:#node[node]
a|
```

```
|physical_protocol
|string
a|The physical network protocol of the FC port.
```

```
|speed
|link:#speed[speed]
a|The physical device speed related properties of the FC port.
```

```
|state
|string
a|The operational state of the FC port.
```

- * startup - The port is booting up.
- * link_not_connected - The port has finished initialization, but a link with the fabric is not established.
- * online - The port is initialized and a link with the fabric has been established.
- * link_disconnected - The link was present at one point on this port but is currently not established.
- * offline_by_user - The port is administratively disabled.
- * offline_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.
- * node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible.

```

|statistics
|link:#statistics[statistics]
a|These are raw performance numbers, such as IOPS latency and throughput.
These numbers are aggregated across all nodes in the cluster and increase
with the uptime of the cluster.

|supported_protocols
|array[string]
a|The network protocols supported by the FC port.

|transceiver
|link:#transceiver[transceiver]
a|Properties of the transceiver connected to the FC port.

|uuid
|string
a|The unique identifier of the FC port.

|wwnn
|string
a|The base world wide node name (WWNN) for the FC port.

|wwpn
|string
a|The base world wide port name (WWPN) for the FC port.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string

```



```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[ID493924ee35377b6eacc3ca6229bd0567]]
```

= Retrieve an FC port

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/network/fc/ports/{uuid}`#
```

Introduced In: 9.6

Retrieves an FC port.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `fabric.name`

* `statistics.+++`

* `metric.+++`

== Related ONTAP commands

* `network fcp adapter show`

== Learn more

* [xref:{relative_path}network_fc_ports_endpoint_overview.html\[DOC /network/fc/ports\]](#)

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```

|path
|True
a|The unique identifier for the FC port.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

```

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|description
|string
a|A description of the FC port.

|enabled
|boolean
a|The administrative state of the FC port. If this property is set to
_false_, all FC connectivity to FC interfaces are blocked. Optional in
PATCH.

|fabric
|link:#fabric[fabric]
a|Properties of the fabric to which the FC port is attached.

|metric
|link:#metric[metric]

```

a|Performance numbers, such as IOPS latency and throughput

|name

|string

a|The FC port name.

|node

|link:#node[node]

a|

|physical_protocol

|string

a|The physical network protocol of the FC port.

|speed

|link:#speed[speed]

a|The physical device speed related properties of the FC port.

|state

|string

a|The operational state of the FC port.

* startup - The port is booting up.

* link_not_connected - The port has finished initialization, but a link with the fabric is not established.

* online - The port is initialized and a link with the fabric has been established.

* link_disconnected - The link was present at one point on this port but is currently not established.

* offlined_by_user - The port is administratively disabled.

* offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.

* node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible.

|statistics

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
|supported_protocols
|array[string]
a|The network protocols supported by the FC port.
```

```
|transceiver
|link:#transceiver[transceiver]
a|Properties of the transceiver connected to the FC port.
```

```
|uuid
|string
a|The unique identifier of the FC port.
```

```
|wwnn
|string
a|The base world wide node name (WWNN) for the FC port.
```

```
|wwpn
|string
a|The base world wide port name (WWPN) for the FC port.
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "description": "Fibre Channel Target Adapter 0a (ACME Fibre Channel
Adapter, rev. 1.0.0, 8G)",
  "fabric": {
    "connected_speed": "16",
    "name": "string",
    "port_address": "52100A",
    "switch_port": "ssan-g620-03:33"
  },
  "metric": {
```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "name": "0a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "physical_protocol": "string",
  "speed": {
    "configured": "auto",
    "maximum": "32"
  },
  "state": "online",
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    }
  },

```

```

    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "supported_protocols": [
    "string"
  ],
  "transceiver": {
    "capabilities": [
      "16"
    ],
    "form-factor": "string",
    "manufacturer": "Acme, Inc.",
    "part_number": "string"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "wwnn": "20:00:00:50:56:b4:13:a8",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]

```


`_links`

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#fabric]
```

```
[.api-collapsible-fifth-title]
```

`fabric`

Properties of the fabric to which the FC port is attached.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|connected
```

```
|boolean
```

`a|`Reports if the physical port has established a connection with the FC fabric.

```
|connected_speed
```

```
|integer
```

`a|`The negotiated data rate between the target FC port and the fabric in gigabits per second.

```
|name
```

```
|string
```

`a|`The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is

explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|port_address

|string

a|The FC port address of the host bus adapter (HBA) physical port.

Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the FC port address given to the physical host bus adapter (HBA) port when the port performs a fabric login (FLOGI).

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a six-digit hexadecimal encoded numeric value.

|switch_port

|string

a|The switch port to which the FC port is connected.

|===

[#iops]

[.api-collapsible-fifth-title]

iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS latency and throughput

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#speed]
[.api-collapsible-fifth-title]
speed
```

The physical device speed related properties of the FC port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|configured
```

```
|string
a|The configured speed of the FC port in gigabits per second.

|maximum
|string
a|The maximum speed supported by the FC port in gigabits per second.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]  
[.api-collapsible-fifth-title]  
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]  
[.api-collapsible-fifth-title]  
throughput_raw
```


Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

```
a|The number of I/O operations observed at the storage object. This should
```

be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw

|link:#latency_raw[latency_raw]

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#transceiver]

[.api-collapsible-fifth-title]

transceiver

Properties of the transceiver connected to the FC port.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|capabilities
|array[integer]
a|The speeds of which the transceiver is capable in gigabits per second.

|form-factor
|string
a|The form factor of the transceiver. Possible values are:

* _sfp_ - Small Form Factor - Pluggable
* _sff_ - Small Form Factor
* _unk_ - Unknown

|manufacturer
|string
a|The manufacturer of the transceiver.

|part_number
|string
a|The part number of the transceiver.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

[[ID5b8d839c846551643dafbdf0f29e90b8]]

= Update an FC port

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-block]#`/network/fc/ports/{uuid}`#

Introduced In: 9.6

Updates an FC port.

== Related ONTAP commands

* `network fcp adapter modify`

== Learn more

* xref:{relative_path}network_fc_ports_endpoint_overview.html[DOC /network/fc/ports]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|The unique identifier for the FC port.

|===

== Request Body

[cols=3*,options=header]

|===

|Name

```

|Type
|Description

|description
|string
a|A description of the FC port.

|enabled
|boolean
a|The administrative state of the FC port. If this property is set to
  _false_, all FC connectivity to FC interfaces are blocked. Optional in
  PATCH.

|fabric
|link:#fabric[fabric]
a|Properties of the fabric to which the FC port is attached.

|name
|string
a|The FC port name.

|node
|link:#node[node]
a|

|physical_protocol
|string
a|The physical network protocol of the FC port.

|speed
|link:#speed[speed]
a|The physical device speed related properties of the FC port.

|state
|string
a|The operational state of the FC port.

* startup - The port is booting up.
* link_not_connected - The port has finished initialization, but a link
  with the fabric is not established.
* online - The port is initialized and a link with the fabric has been

```

established.

* link_disconnected - The link was present at one point on this port but is currently not established.

* offlined_by_user - The port is administratively disabled.

* offlined_by_system - The port is set to offline by the system. This happens when the port encounters too many errors.

* node_offline - The state information for the port cannot be retrieved. The node is offline or inaccessible.

|supported_protocols

|array[string]

a|The network protocols supported by the FC port.

|transceiver

|link:#transceiver[transceiver]

a|Properties of the transceiver connected to the FC port.

|uuid

|string

a|The unique identifier of the FC port.

|wwnn

|string

a|The base world wide node name (WWNN) for the FC port.

|wwpn

|string

a|The base world wide port name (WWPN) for the FC port.

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "description": "Fibre Channel Target Adapter 0a (ACME Fibre Channel Adapter, rev. 1.0.0, 8G)",

 "fabric": {

 "connected_speed": "16",

```

    "name": "string",
    "port_address": "52100A",
    "switch_port": "ssan-g620-03:33"
  },
  "name": "0a",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "physical_protocol": "string",
  "speed": {
    "configured": "auto",
    "maximum": "32"
  },
  "state": "online",
  "supported_protocols": [
    "string"
  ],
  "transceiver": {
    "capabilities": [
      "16"
    ],
    "form-factor": "string",
    "manufacturer": "Acme, Inc.",
    "part_number": "string"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "wwnn": "20:00:00:50:56:b4:13:a8",
  "wwpn": "20:00:00:50:56:b4:13:a8"
}
====

== Response

```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

```



```
| 5374085
| The node where the Fibre Channel port is located is offline.

| 5374087
| The Fibre Channel port has active Fibre Channel interfaces and cannot be
disabled.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#fabric]
[.api-collapsible-fifth-title]
fabric
```

Properties of the fabric to which the FC port is attached.

```
[cols=3*,options=header]
|===
|Name
|Type
```

|Description

|connected

|boolean

a|Reports if the physical port has established a connection with the FC fabric.

|connected_speed

|integer

a|The negotiated data rate between the target FC port and the fabric in gigabits per second.

|name

|string

a|The name of the fabric to which the port is connected. This is only available when the FC port is connected to a fabric.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|port_address

|string

a|The FC port address of the host bus adapter (HBA) physical port.

Each FC port in an FC switched fabric has its own unique FC port address for routing purposes. The FC port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the FC port address given to the physical host bus adapter (HBA) port when the port performs a fabric login (FLOGI).

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a six-digit hexadecimal encoded numeric value.

|switch_port

|string

a|The switch port to which the FC port is connected.

|===

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

|===

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput]
```

```
[.api-collapsible-fifth-title]
```

throughput

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS latency and throughput

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
```

of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#node]

[.api-collapsible-fifth-title]

node

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|

|uuid

|string

a|

|===

[#speed]

```
[.api-collapsible-fifth-title]
```

```
speed
```

The physical device speed related properties of the FC port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|configured
```

```
|string
```

```
a|The configured speed of the FC port in gigabits per second.
```

```
|maximum
```

```
|string
```

```
a|The maximum speed supported by the FC port in gigabits per second.
```

```
|===
```

```
[#iops_raw]
```

```
[.api-collapsible-fifth-title]
```

```
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```



```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw performance numbers, such as IOPS latency and throughput.

These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
```

```
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp
```

```

|string
a|The timestamp of the performance data.

|===

[#transceiver]
[.api-collapsible-fifth-title]
transceiver

Properties of the transceiver connected to the FC port.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|capabilities
|array[integer]
a|The speeds of which the transceiver is capable in gigabits per second.

|form-factor
|string
a|The form factor of the transceiver. Possible values are:

* _sfp_ - Small Form Factor - Pluggable
* _sff_ - Small Form Factor
* _unk_ - Unknown

|manufacturer
|string
a|The manufacturer of the transceiver.

|part_number
|string
a|The part number of the transceiver.

|===

```

```
[#fc_port]
[.api-collapsible-fifth-title]
fc_port
```

A Fibre Channel (FC) port is the physical port of an FC adapter on an ONTAP cluster node that can be connected to an FC network to provide FC network connectivity. An FC port defines the location of an FC interface within the ONTAP cluster.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|description
|string
a|A description of the FC port.
```

```
|enabled
|boolean
a|The administrative state of the FC port. If this property is set to
  _false_, all FC connectivity to FC interfaces are blocked. Optional in
  PATCH.
```

```
|fabric
|link:#fabric[fabric]
a|Properties of the fabric to which the FC port is attached.
```

```
|name
|string
a|The FC port name.
```

```
|node
|link:#node[node]
a|
```

```
|physical_protocol
|string
a|The physical network protocol of the FC port.
```

```

|speed
|link:#speed[speed]
a|The physical device speed related properties of the FC port.

|state
|string
a|The operational state of the FC port.

* startup - The port is booting up.
* link_not_connected - The port has finished initialization, but a link
with the fabric is not established.
* online - The port is initialized and a link with the fabric has been
established.
* link_disconnected - The link was present at one point on this port but
is currently not established.
* offlined_by_user - The port is administratively disabled.
* offlined_by_system - The port is set to offline by the system. This
happens when the port encounters too many errors.
* node_offline - The state information for the port cannot be retrieved.
The node is offline or inaccessible.

|supported_protocols
|array[string]
a|The network protocols supported by the FC port.

|transceiver
|link:#transceiver[transceiver]
a|Properties of the transceiver connected to the FC port.

|uuid
|string
a|The unique identifier of the FC port.

|wwnn
|string
a|The base world wide node name (WWNN) for the FC port.

|wwpn
|string
a|The base world wide port name (WWPN) for the FC port.

```

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

|===

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDc3d570dfac058f79b49e42086108c65a]]
= Retrieve FC port historical performance metrics

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/fc/ports/{uuid}/metrics`#

*Introduced In:* 9.8

Retrieves historical performance metrics for an FC port

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|duration
|string
|query
|False
a|Filter by duration

```

```
|timestamp  
|string  
|query  
|False  
a|Filter by timestamp
```

```
|throughput.write  
|integer  
|query  
|False  
a|Filter by throughput.write
```

```
|throughput.read  
|integer  
|query  
|False  
a|Filter by throughput.read
```

```
|throughput.total  
|integer  
|query  
|False  
a|Filter by throughput.total
```

```
|latency.total  
|integer  
|query  
|False  
a|Filter by latency.total
```

```
|latency.read  
|integer  
|query  
|False  
a|Filter by latency.read
```

```
|latency.other  
|integer  
|query  
|False  
a|Filter by latency.other
```



```
|latency.write
|integer
|query
|False
a|Filter by latency.write
```

```
|status
|string
|query
|False
a|Filter by status
```

```
|iops.total
|integer
|query
|False
a|Filter by iops.total
```

```
|iops.read
|integer
|query
|False
a|Filter by iops.read
```

```
|iops.other
|integer
|query
|False
a|Filter by iops.other
```

```
|iops.write
|integer
|query
|False
a|Filter by iops.write
```

```
|uuid
|string
|path
```

```

|True
a|Unique identifier of the FC port.

|interval
|string
|query
|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:

* 1h: Metrics over the most recent hour sampled over 15 seconds.
* 1d: Metrics over the most recent day sampled over 5 minutes.
* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

```

```

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#records[records]]
a|

|===

```

```

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
```

Name
Type
Description
other
integer
a Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read
integer
a Performance metric for read I/O operations.
total
integer
a Performance metric aggregated over all types of I/O operations.
write
integer
a Performance metric for write I/O operations.
===
[#throughput]
[.api-collapsible-fifth-title]
throughput
The rate of throughput bytes per second observed at the storage object.
[cols=3*,options=header]
===
Name
Type
Description
read
integer
a Performance metric for read I/O operations.
total


```

|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#records]
[.api-collapsible-fifth-title]
records

Performance numbers, such as IOPS latency and throughput

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status

```

```
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage HTTP proxy configuration

:leveloffset: +1

```

```
[[IDefe094203686f8f886a8a1faa280493b]]  
= Network http-proxy endpoint overview
```

```
== Overview
```

Configuration of an HTTP proxy for an SVM or a Cluster IPspace.

```
== Retrieve HTTP proxy information
```

The HTTP proxy GET operation retrieves all configurations for an SVM or a Cluster IPspace via `'/api/cluster'`.

```
== Examples
```

```
=== Retrieving all fields for all HTTP proxy configurations
```

```
----
```

```
# The API:  
/api/network/http-proxy
```

```
# The call:  
curl -X GET "https://<mgmt-ip>/api/network/http-  
proxy?fields=*&return_records=true&return_timeout=15" -H "accept:  
application/json"
```

```
# The response:  
{  
  "records": [  
    {  
      "uuid": "4133a1fc-7228-11e9-b40c-005056bb4f0c",  
      "svm": {  
        "name": "vs1",  
        "uuid": "4133a1fc-7228-11e9-b40c-005056bb4f0c"  
      },  
      "server": "server1.example.com",  
      "port": 3128  
    },  
    {  
      "uuid": "96219ce3-7214-11e9-828c-005056bb4f0c",  
      "svm": {  
        "name": "cluster-1",  
        "uuid": "96219ce3-7214-11e9-828c-005056bb4f0c"  
      }  
    }  
  ]  
}
```

```

    },
    "ipospace": {
        "uuid": "7433520f-7214-11e9-828c-005056bb4f0c",
        "name": "Default"
    },
    "server": "1.1.1.",
    "port": 3128
}
],
"num_records": 2
}
-----

```

=== Retrieving the HTTP proxy configuration for a specific SVM

```

-----

# The API:
/api/network/http-proxy/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/http-proxy/96219ce3-7214-11e9-828c-005056bb4f0c" -H "accept: application/json"

# The response
{
  "uuid": "96219ce3-7214-11e9-828c-005056bb4f0c",
  "svm": {
    "name": "cluster-1",
    "uuid": "96219ce3-7214-11e9-828c-005056bb4f0c"
  },
  "ipospace": {
    "uuid": "7433520f-7214-11e9-828c-005056bb4f0c",
    "name": "Default"
  },
  "server": "1.1.1.1",
  "port": 3128
}
-----

```

=== Creating an HTTP proxy configuration

You can use the HTTP proxy POST operation to create an HTTP proxy configuration for the specified SVM.

=== Examples

=== Creating an HTTP proxy configuration for a particular SVM

The API:

/api/network/http-proxy

The call:

```
curl -X POST "https://<mgmt-ip>/api/network/http-proxy" -H "accept:
application/json" -H "Content-Type: application/json" -d "{  \"port\":
3128,  \"server\": \"1.1.1.1\",  \"svm\": {      \"name\": \"cluster-1\"
}}"
```

=== Creating an HTTP proxy configuration for a particular IPspace

The API:

/api/network/http-proxy

The call:

```
curl -X POST "https://<mgmt-ip>/api/network/http-proxy" -H "accept:
application/json" -H "Content-Type: application/json" -d "{  \"ipspace\":
{      \"name\": \"Default\"  },  \"port\": 3128,  \"server\": \"1.1.1.1\"}"
```

== Update an HTTP proxy configuration for a specified SVM

You can use the HTTP proxy PATCH operation to update the HTTP proxy configuration for the specified SVM.

== Example

The following example shows how a PATCH operation is used to update an HTTP proxy configuration for a specific SVM:

The API:

/api/network/http-proxy/{uuid}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/network/http-proxy/96219ce3-7214-
11e9-828c-005056bb4f0c" -H "accept: application/json" -H "Content-Type:
application/json" -d "{      \"port\": 3128,  \"server\":
\"server2.example.com\"}"
```

====

== Delete an HTTP proxy configuration for a specified SVM

You can use the HTTP proxy DELETE operation to delete the HTTP proxy configuration for the specified SVM.

== Example

The following example shows how a DELETE operation is used to delete an HTTP proxy configuration for a specific SVM:

```
# The API:
/api/network/http-proxy/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/http-proxy/96219ce3-7214-11e9-828c-005056bb4f0c" -H "accept: application/json"
```

```
[[ID81e7eebead808474b151781f3c101062]]
= Retrieve HTTP proxy configurations for all SVMs and cluster IPspaces
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/http-proxy`#
```

Introduced In: 9.7

Retrieves the HTTP proxy configurations of all the SVMs and Cluster IPspaces.

== Related ONTAP commands

* `vserver http-proxy show`

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description

|server
|string
|query
|False
a|Filter by server

|port
|integer
|query
|False
a|Filter by port

|scope
|string
|query
|False
a|Filter by scope

|ipspace.name
|string
|query
|False
a|Filter by ipspace.name

|ipspace.uuid
|string
|query
|False
a|Filter by ipspace.uuid

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```



```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

```

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of HTTP proxy records

|records
|array[link:#network_http_proxy[network_http_proxy]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{

```

```

    "_links": {
      "next": {
        "href": "/api/resourcelink"
      },
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "records": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ipospace": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "exchange",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "port": "3128",
        "scope": "string",
        "server": "string",
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "uuid": "string"
      }
    ]
  }
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

```

```
[#ipspace]
[.api-collapsible-fifth-title]
ipspace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|IPspace name
```

```
|uuid
|string
a|IPspace UUID
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
```

```

|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#network_http_proxy]
[.api-collapsible-fifth-title]
network_http_proxy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
is supplied on input.

|port
|integer
a|The port number on which the HTTP proxy service is configured on the
proxy server.

|scope
|string
a|Set to "`svm`" for proxy owned by an SVM. Otherwise, set to "cluster".

|server
|string
a|The fully qualified domain name (FQDN) or IP address of the proxy
server.

```

```
|svm
|link:#svm[svm]
a|

|uuid
|string
a|The UUID that uniquely identifies the HTTP proxy.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```



```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[ID27fbe88d795b194bd1bc529d8614fe55]]
= Create an HTTP proxy configuration for an SVM or cluster IPspace

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/network/http-proxy`#
```

***Introduced In:* 9.7**

Creates an HTTP proxy configuration for an SVM or a Cluster IPspace.
Important notes:

- * IPv6 must be enabled if IPv6 family addresses are specified in the "server" field.
- * The server and the port combination specified using the "server" and "port" fields is validated during this operation. The validation will fail in the following scenarios:

- *** The HTTP proxy service is not configured on the server.

```

*** The HTTP proxy service is not running on the specified port.

*** The server is unreachable.

== Required properties

* SVM-scoped HTTP proxy

*** `svm.uuid` or `svm.name` - Existing SVM in which to create the HTTP
proxy.

* Cluster-scoped HTTP proxy

*** `ipspace.uuid` or `ipspace.name` - Existing Cluster IPspace in which
to create the HTTP proxy.

* `server` - HTTP proxy server FQDN or IP address.
* `port` - HTTP proxy server port.

== Related ONTAP commands

* `vserver http-proxy create`

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ipSPACE
|link:#ipSPACE[ipSPACE]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
is supplied on input.

|port
|integer
a|The port number on which the HTTP proxy service is configured on the
proxy server.

|scope
|string
a|Set to "`svm`" for proxy owned by an SVM. Otherwise, set to "cluster".

|server
|string
a|The fully qualified domain name (FQDN) or IP address of the proxy
server.

|svm
|link:#svm[svm]
a|

|uuid
|string
a|The UUID that uniquely identifies the HTTP proxy.

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]

```

```
{
  "ipospace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "port": "3128",
  "scope": "string",
  "server": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
====

== Response
```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of HTTP proxy records


|records
|array[link:#network_http_proxy[network_http_proxy]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "ipspace": {
        "name": "exchange",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": "3128",
      "scope": "string",
      "server": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "string"
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 26214473

| HTTP proxy configuration is not valid.

| 26214476

| The "IPspace" parameter should not be specified in the SVM context.

| 26214477

| The specified IPspace does not exist.

| 23724130

| Cannot use an IPv6 name server address because there are no IPv6 interfaces.

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

=====

[#href]

[.api-collapsible-fifth-title]

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

|string

a|

|===

[#_links]

[.api-collapsible-fifth-title]

```
_links
[#ipspace]
[.api-collapsible-fifth-title]
ipspace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|IPspace name
```

```
|uuid
|string
a|IPspace UUID
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

|===

```
[#network_http_proxy]
[.api-collapsible-fifth-title]
network_http_proxy
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|ipSPACE
|link:#ipSPACE[ipSPACE]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
is supplied on input.
```

```
|port
|integer
a|The port number on which the HTTP proxy service is configured on the
proxy server.
```

```
|scope
|string
a|Set to "`svm`" for proxy owned by an SVM. Otherwise, set to "cluster".
```

```
|server
|string
a|The fully qualified domain name (FQDN) or IP address of the proxy
server.
```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|The UUID that uniquely identifies the HTTP proxy.
```

|===


```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

:leveloffset: -1

[[IDcfe51f0906bb456967352b24a4d2961a]]

= Deletes the http proxy configuration of the specified svm or cluster
ipspace

related ontap commands

* `vserver http-proxy delete`

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/network/http-proxy/{uuid}`#

Introduced In: 9.7

Deletes the HTTP proxy configuration of the specified SVM or Cluster
IPspace.

== Related ONTAP commands

* `vserver http-proxy delete`

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

```
|Required
|Description

|uuid
|string
|path
|True
a|HTTP proxy UUID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
```

```

    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

```

== Definitions

```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

```

```

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|arguments

```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[ID7e3e5eb00369048a566a10f88abf1a42]]
```

```
= Displays the http proxy server, port, and ipspace of the specified svm  
or cluster ipspace
```

```
### related ontap commands
```

```
* `vserver http-proxy show`
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/network/http-proxy/{uuid}`#
```

```
*Introduced In:* 9.7
```

```
Displays the HTTP proxy server, port, and IPspace of the specified SVM or  
Cluster IPspace.
```

```
== Related ONTAP commands
```

```
* `vserver http-proxy show`
```

```

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|HTTP proxy UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ipSPACE
|link:#ipSPACE[ipSPACE]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
is supplied on input.

```

```

|port
|integer
a|The port number on which the HTTP proxy service is configured on the
proxy server.

|scope
|string
a|Set to "`svm`" for proxy owned by an SVM. Otherwise, set to "cluster".

|server
|string
a|The fully qualified domain name (FQDN) or IP address of the proxy
server.

|svm
|link:#svm[svm]
a|

|uuid
|string
a|The UUID that uniquely identifies the HTTP proxy.

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}

```

```

},
"port": "3128",
"scope": "string",
"server": "string",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "string"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ]
  }
}

```



```

    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

```

```
[#ipspace]
[.api-collapsible-fifth-title]
ipspace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|IPspace name
```

```
|uuid
|string
a|IPspace UUID
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[IDc570910e3a2141e0c763a54a8358336a]]
```

```
= Updates the proxy server, port, username, and password parameters  
important notes:
```

```
* ipv6 must be enabled if ipv6 family addresses are specified in the  
"server" field
```

```
* the server and the port combination specified using the "server" and  
"port" fields is validated during this operation the validation will fail  
in the following scenarios:
```

```
    * the http proxy service is not configured on the server
```

```
    * the http proxy service is not running on the specified port
```

```
    * the server is unreachable
```

```
### related ontap commands
```

```
* `vserver http-proxy modify`
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-  
block]#`/network/http-proxy/{uuid}`#
```

```
*Introduced In:* 9.7
```

Updates the proxy server, port, username, and password parameters.

Important notes:

- * IPv6 must be enabled if IPv6 family addresses are specified in the "server" field.

- * The server and the port combination specified using the "server" and "port" fields is validated during this operation. The validation will fail in the following scenarios:

- *** The HTTP proxy service is not configured on the server.

- *** The HTTP proxy service is not running on the specified port.

- *** The server is unreachable.

== Related ONTAP commands

- * ``vserver http-proxy modify``

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|HTTP proxy UUID

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

```

|Description

|port
|integer
a|The port number on which the HTTP proxy service is configured on the
proxy server.

|scope
|string
a|Set to "`svm`" for proxy owned by an SVM. Otherwise, set to "cluster".

|server
|string
a|The fully qualified domain name (FQDN) or IP address of the proxy
server.

|uuid
|string
a|The UUID that uniquely identifies the HTTP proxy.

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "port": "3128",
  "scope": "string",
  "server": "string",
  "uuid": "string"
}
====

== Response

```

Status: 200, Ok

```
== Error
```

ONTAP Error Response Codes

|===

| Error Code | Description

| 26214473

| The HTTP proxy configuration is not valid.

| 23724130

| Cannot use an IPv6 name server address because there are no IPv6 interfaces.

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

=====

[#href]

[.api-collapsible-fifth-title]

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

|string

a|

|===

[#_links]

[.api-collapsible-fifth-title]

_links

[#ipspace]

[.api-collapsible-fifth-title]

ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|IPspace name
```

```
|uuid
```

```
|string
```

```
a|IPspace UUID
```

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the SVM.
```

```
|===
```



```

[#network_http_proxy]
[.api-collapsible-fifth-title]
network_http_proxy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|port
|integer
a|The port number on which the HTTP proxy service is configured on the
proxy server.

|scope
|string
a|Set to "`svm`" for proxy owned by an SVM. Otherwise, set to "cluster".

|server
|string
a|The fully qualified domain name (FQDN) or IP address of the proxy
server.

|uuid
|string
a|The UUID that uniquely identifies the HTTP proxy.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string

```

```

a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

= Manage BGP peer groups

:leveloffset: +1

[[ID5a78070681a8c9893717d7657e677fd6]]

= Network IP BGP peer-groups endpoint overview

== Overview

The following operations are supported:

- * Creation: POST network/ip/bgp/peer-groups
- * Collection Get: GET network/ip/bgp/peer-groups
- * Instance Get: GET network/ip/bgp/peer-groups/{uuid}
- * Instance Patch: PATCH network/ip/bgp/peer-groups/{uuid}
- * Instance Delete: DELETE network/ip/bgp/peer-groups/{uuid}

== Retrieving network BGP sessions information

The IP BGP peer-groups GET API retrieves and displays relevant information pertaining to the BGP peer-groups configured in the cluster. The response can contain a list of multiple BGP peer-groups or a specific peer-group. Each BGP peer-group represents a BGP session configured between a local interface and a peer router.

== Examples

=== Retrieving all BGP peer-groups in the cluster

The following example shows the list of all BGP peer-groups configured in a cluster.

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ip/bgp/peer-groups
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/network/ip/bgp/peer-groups" -H "accept: application/hal+json"
```

```
# The response:
{
  "records": [
    {
      "uuid": "5f22ae9d-87b2-11e9-a3a6-005056bb81a4",
      "name": "pg1",
      "_links": {
        "self": {
          "href": "/api/network/ip/bgp/peer-groups/5f22ae9d-87b2-11e9-a3a6-005056bb81a4"
        }
      }
    },
    {
      "uuid": "5fd08be3-87b2-11e9-952f-005056bb2170",
      "name": "pg2",
      "_links": {
        "self": {
          "href": "/api/network/ip/bgp/peer-groups/5fd08be3-87b2-11e9-952f-005056bb2170"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/network/ip/bgp/peer-groups"
    }
  }
}
----
```

'''

=== Retrieving a specific BGP peer-group

The following example shows the response when a specific BGP peer-group is requested. The system returns an error when there is no peer-group with the requested UUID.

'''

```

# The API:
/api/network/ip/bgp/peer-groups/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/bgp/peer-groups/5fd08be3-87b2-11e9-952f-005056bb2170" -H "accept: application/hal+json"

# The response:
{
  "uuid": "5fd08be3-87b2-11e9-952f-005056bb2170",
  "name": "pg2",
  "ipspace": {
    "uuid": "84fd3375-879a-11e9-a3a6-005056bb81a4",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/84fd3375-879a-11e9-a3a6-005056bb81a4"
      }
    }
  },
  "local": {
    "interface": {
      "uuid": "5e76a305-87b2-11e9-952f-005056bb2170",
      "name": "bgp2",
      "ip": {
        "address": "10.10.10.2"
      }
    },
    "port": {
      "uuid": "f8ff73de-879a-11e9-952f-005056bb2170",
      "name": "e0h",
      "node": {
        "name": "node1"
      }
    }
  },
  "peer": {
    "address": "10.10.10.1",
    "asn": 65501
  },
  "state": "up",
  "_links": {
    "self": {
      "href": "/api/network/ip/bgp/peer-groups/5fd08be3-87b2-11e9-952f-005056bb2170"
    }
  }
}

```

```

}
}
----

'''

=== Retrieving specific fields and limiting the output using filters

The following example shows the response when a filter is applied
(location.port.node.name=node1) and only certain fields are requested.
Filtered fields are in the output in addition to the default fields and
requested fields.

'''

----

# The API:
/api/network/ip/bgp/peer-groups

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/bgp/peer-
groups?local.port.node.name=node1&fields=local.interface.ip,peer" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "5f22ae9d-87b2-11e9-a3a6-005056bb81a4",
      "name": "pg1",
      "local": {
        "interface": {
          "ip": {
            "address": "10.10.10.1"
          }
        },
        "port": {
          "node": {
            "name": "node1"
          }
        }
      },
      "peer": {
        "address": "10.10.10.2",
        "asn": 65501
      }
    }
  ]
}

```

```

    },
    "_links": {
      "self": {
        "href": "/api/network/ip/bgp/peer-groups/5f22ae9d-87b2-11e9-a3a6-005056bb81a4"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ip/bgp/peer-groups?local.port.node.name=node1&fields=local.interface.ip,peer"
  }
}
}
}
----
'''

```

== Creating a BGP peer-group

The BGP peer-group POST API is used to create a peer-group as shown in the following examples.

```
'''
```

== Examples

=== Creating a BGP peer-group with an existing interface

The following example shows how to create a BGP peer-group between an existing interface "bgp1" and peer router with the address "10.10.10.10". The local interface "bgp1" needs to support the management-bgp service, otherwise the system returns an error.

```
'''
```

```
----
```

The API:

/api/network/ip/bgp/peer-groups

The call:

```
curl -X POST "https://<mgmt-ip>/api/network/ip/bgp/peer-groups?return_records=true" -d'{"name": "newPg", "ipspace.name": "Default",
```

```
"local.interface.name": "bgp1", "peer.address": "10.10.10.10"}'
```

```
# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "e3faacc6-87cb-11e9-a3a6-005056bb81a4",
      "name": "newPg",
      "ipspace": {
        "name": "Default"
      },
      "local": {
        "interface": {
          "name": "bgp1"
        }
      },
      "peer": {
        "address": "10.10.10.10"
      },
      "_links": {
        "self": {
          "href": "/api/network/ip/bgp/peer-groups/e3faacc6-87cb-11e9-a3a6-005056bb81a4"
        }
      }
    }
  ]
}
----
```

```
'''
```

=== Creating a BGP peer-group and provisioning a new local interface

The following example shows how to create a BGP peer-group with any local interface. If the local interface doesn't exist, the system will create it first before creating the peer-group.

```
'''
----
```

```
# The API:
/api/network/ip/bgp/peer-groups
```



```

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/bgp/peer-
groups?return_records=true" -d'{"name": "newPg1",
"ipspace.name":"Default", "local": {"interface": {"name": "newlif"}, "ip":
{"address": "9.9.9.9", "netmask": "24"}, "port": {"name": "e0f", "node":
{"name": "node1"}}}, "peer.address":"10.10.10.10"}'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "c292f069-8872-11e9-a3a6-005056bb81a4",
      "name": "newPg1",
      "ipspace": {
        "name": "Default"
      },
      "local": {
        "interface": {
          "name": "newlif"
        },
        "port": {
          "name": "e0f",
          "node": {
            "name": "node1"
          }
        }
      },
      "peer": {
        "address": "10.10.10.10"
      },
      "_links": {
        "self": {
          "href": "/api/network/ip/bgp/peer-groups/c292f069-8872-11e9-a3a6-
005056bb81a4"
        }
      }
    }
  ]
}
----

'''

== Updating BGP peer-groups

```

The BGP peer-groups PATCH API is used to update attributes of a peer-group.

'''

== Examples

=== Updating the peer router address

The following example shows how the PATCH request changes the peer router IP address.

'''

The API:

/api/network/ip/bgp/peer-groups/{uuid}

The call:

curl -X PATCH "https://<mgmt-ip>/api/network/ip/bgp/peer-groups/80d271c9-1f43-11e9-803e-005056a7646a" -H "accept: application/hal+json" -d

'{"peer.address": "10.10.10.20" }'

{

}

'''

=== Updating the peer-group to a new name

The following example shows how the PATCH request renames the peer-group.

'''

The API:

/api/network/ip/bgp/peer-groups/{uuid}

The call:

curl -X PATCH "https://<mgmt-ip>/api/network/ip/bgp/peer-groups/80d271c9-1f43-11e9-803e-005056a7646a" -H "accept: application/hal+json" -d

'{"name": "NewName" }'

{

}

'''

== Deleting BGP peer-groups

The BGP peer-groups DELETE API is used to delete an BGP peer-group.

'''

== Example

=== Deleting a BGP peer-group

The following DELETE request deletes a BGP peer-group.

'''

The API:

/api/network/ip/bgp/peer-group/{uuid}

The call:

curl -X DELETE "https://<mgmt-ip>/api/network/ip/bgp/peer-groups/80d271c9-1f43-11e9-803e-005056a7646a"

{
}

'''

[[IDf1656caalca022a06ba40b77d854bd54]]

= Retrieve all BGP peer group details for VIP

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/network/ip/bgp/peer-groups`#

Introduced In: 9.7

Retrieves the details of all BGP peer groups for VIP.

== Related ONTAP Commands

```
* `network bgp peer-group show`
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|state
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by state
```

```
|local.port.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by local.port.name
```

```
|local.port.node.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by local.port.node.name
```

```
|local.port.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by local.port.uuid
```

```
|local.interface.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by local.interface.name
```

```
|local.interface.uuid  
|string  
|query  
|False  
a|Filter by local.interface.uuid
```

```
|local.interface.ip.address  
|string  
|query  
|False  
a|Filter by local.interface.ip.address
```

```
|uuid  
|string  
|query  
|False  
a|Filter by uuid
```

```
|peer.asn  
|integer  
|query  
|False  
a|Filter by peer.asn
```

```
|peer.address  
|string  
|query  
|False  
a|Filter by peer.address
```

```
|name  
|string  
|query  
|False  
a|Filter by name
```

```
|ipspace.name  
|string  
|query
```

```

|False
a|Filter by ipspace.name

|ipspace.uuid
|string
|query
|False
a|Filter by ipspace.uuid

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|

|records
|array[link:#bgp_peer_group[bgp_peer_group]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
```

```

    "href": "/api/resourcelink"
  }
},
"num_records": "1",
"records": [
  {
    "ipospace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "exchange",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "local": {
      "interface": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "ip": {
          "address": "10.10.10.7"
        },
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "port": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "elb",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    },
    "name": "bgpv4peer",
    "peer": {
      "address": "10.10.10.7"
    },
    "state": "string",

```



```

      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions

```

```

[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#ipspace]
[.api-collapsible-fifth-title]
ipspace
```

Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|IPspace name
```

```
|uuid
|string
a|IPspace UUID
```

```
|===
```

```
[#ip]
[.api-collapsible-fifth-title]
ip
```

IP information

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|===

[#interface]
[.api-collapsible-fifth-title]
interface

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ip
|link:#ip[ip]
a|IP information

|name
|string
a|The name of the interface.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|===

```

```

[#ip]
[.api-collapsible-fifth-title]
ip

IP information to create a new interface.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.


|===


[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of node on which the port is located.


|===


[#port]

```

```
[.api-collapsible-fifth-title]
```

```
port
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#local]
```

```
[.api-collapsible-fifth-title]
```

```
local
```

Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|interface
```

```
|link:#interface[interface]
```

```
a|
```

```
|port
|link:#port[port]
a|
```

```
|===
```

```
[#peer]
[.api-collapsible-fifth-title]
peer
```

Information describing the router to peer with

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|address
|string
a|Peer router address
```

```
|asn
|integer
a|Autonomous system number of peer
```

```
|===
```

```
[#bgp_peer_group]
[.api-collapsible-fifth-title]
bgp_peer_group
```

A BGP peer group between a local network interface and a router, for the purpose of announcing VIP interface locations for SVMs in this IPspace.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```

|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name is supplied on input.

|local
|link:#local[local]
a|Information describing the local interface that is being used to peer
with a router using BGP. On a POST operation, an existing BGP interface is
used by specifying the interface, or create a new one by specifying the
port and IP address.

|name
|string
a|Name of the peer group

|peer
|link:#peer[peer]
a|Information describing the router to peer with

|state
|string
a|State of the peer group

|uuid
|string
a|UUID of the peer group

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

```

====

[[ID571c1866b83395fb505a7e7454ce1b82]]

= Create a new BGP peer group for VIP

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/network/ip/bgp/peer-groups`#

Introduced In: 9.7

Creates a new BGP peer group for VIP. Multipath-routing is turned on cluster-wide automatically if the peer group being created results in multiple paths being available for an existing or future VIP interface.

== Required properties

* `name` - Name of the peer-group to create.

* `ipspace.name` or `ipspace.uuid`

*** Required with local.interface.name to identify a local interface

*** Optional when local.interface.uuid is specified

* `local.interface.uuid` or `local.interface.name`

*** Required when specifying an existing local interface.

* `local.interface.name`, `local.ip` and `local.port`

*** Required to create a new local interface.

* `peer.address` - IP address of the peer router

== Related ONTAP commands

* `network bgp peer-group create`

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

```

|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name is supplied on input.

|local
|link:#local[local]
a|Information describing the local interface that is being used to peer
with a router using BGP. On a POST operation, an existing BGP interface is
used by specifying the interface, or create a new one by specifying the
port and IP address.

|name
|string
a|Name of the peer group

|peer
|link:#peer[peer]
a|Information describing the router to peer with

```

```
|state
|string
a|State of the peer group
```

```
|uuid
|string
a|UUID of the peer group
```

```
|===
```

```
.Example request
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "ipspace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local": {
    "interface": {
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "ip": {
      "address": "10.10.10.7",
      "netmask": "24"
    },
    "port": {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "bgpv4peer",
  "peer": {
    "address": "10.10.10.7"
  },
  "state": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

```
====
```

```
== Response
```

Status: 201, Created

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 1376963
```

```
| Duplicate IP address is specified.
```

```
| 1966133
```

```
| Since masking an address with a netmask represents an entire IP subnet,  
the masked and unmasked IP addresses cannot be the same.
```

```
| 1966267
```

```
| IPv6 addresses must have a prefix length of 64.
```

```
| 1966269
```

```
| IPv4 addresses must have a netmask length between 1 and 32.
```

```
| 1967082
```

```
| IPspace name and UUID must match if both are given.
```

```
| 1967155
```

```
| The specified local.port.name does not match the location.port.name for  
the specified local.interface.
```

```
| 1967156
```

```
| The specified local.port.node.name does not match the  
location.port.node.name for the specified local.interface.
```

```
| 1967157
```

```
| The specified local.port.uuid does not match the location.port.uuid for  
the specified local.interface.
```

```
| 1967158
```

```
| The specified local.interface.name does not exist in the associated
```

IPspace. local.ip.address and local.ip.netmask are required to create a new LIF.

| 1967159

| local.interface does not support management-bgp service.

| 1967160

| The specified local.interface.name does not match the specified interface name of local.interface.uuid.

| 1967161

| The specified local.interface.uuid does not exist in the specified IPspace.

| 1967162

| Either local.interface or local.ip and local.port are required to specify a local LIF.

| 1967163

| The specified local.port.name does not match the specified port name of local.port.uuid.

| 1967164

| The specified local.port.node.name does not match the specified node name of local.port.uuid.

| 1967165

| The specified local.port does not exist.

| 1967166

| ipspace.uuid or ipspace.name must be provided with local.interface.name together to identify a LIF.

| 1967167

| Internal error. Failed to update BGP configuration for node. Retry the command, if necessary.

| 1967168

| Internal error. Failed to create a VIP port for IPspace on node. Retry the command, if necessary.

| 1967169

| Internal error. BGP configuration changed during the operation. Retry the command, if necessary.

| 1967170

| Internal error. VIP port configuration changed during the operation.

Retry the command, if necessary.

| 1967171

| Internal error. Fail to access or update BGP peer group. Retry the command, if necessary.

| 1967172

| Peer group could not be updated because IPspace does not exist. Retry the command, if necessary.

| 1967173

| The specified local.ip.address does not match the address for the specified local.interface.

| 1967174

| The specified local.ip.netmask does not match the netmask for the specified local.interface.

| 1967176

| The specified local.interface.name does not exist in the associated IPspace. local.port.name, local.port.node.name, or local.port.uuid is required to create a new LIF.

| 1967177

| Internal error. Failed to access the local interface. Retry the command, if necessary.

| 1967178

| The IPv6 address specified with local.ip.address is not supported because it is link-local, multicast, v4-compatible, v4-mapped, loopback or "::".

| 1967179

| The IPv4 address specified with local.ip.address is not supported because it is multicast, loopback or 0.0.0.0.

| 53281985

| Internal error. Failed to update BGP peer group because BGP LIF moved during the operation. Wait a few minutes and try the command again.

| 53282006

| BGP peer group could not be updated to use a peer address because the value provided is not a valid peer address. If necessary, try the command again with a routable host address.

| 53282007

| BGP peer group could not be updated to use a peer address because the

address represents a different address family to the address of the associated BGP LIF. If necessary, try the command again with a matching address family.

```
| 53282018
| Failed to create BGP peer group because an existing peer group has
| already established a BGP session between LIF and peer address. If
| necessary, try the command again with a different BGP LIF or a different
| peer address.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

====

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

|===

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#ipospace]
[.api-collapsible-fifth-title]
ipospace
```

Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|name
|string
a|IPspace name


|uuid
|string
a|IPspace UUID


|===


[#ip]
[.api-collapsible-fifth-title]
ip

IP information


[#interface]
[.api-collapsible-fifth-title]
interface


[cols=3*,options=header]
|===
|Name
|Type
|Description

|ip
|link:#ip[ip]
a|IP information


|name
|string
a|The name of the interface.


|uuid
|string
a|The UUID that uniquely identifies the interface.

```

```
|===
```

```
[#ip]
```

```
[.api-collapsible-fifth-title]
```

```
ip
```

IP information to create a new interface.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|address
```

```
|string
```

```
a|IPv4 or IPv6 address
```

```
|netmask
```

```
|string
```

```
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.
```

```
|===
```

```
[#node]
```

```
[.api-collapsible-fifth-title]
```

```
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|Name of node on which the port is located.
```

```
|===
```

```
[#port]  
[.api-collapsible-fifth-title]  
port
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|
```

```
|node  
|link:#node[node]  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#local]  
[.api-collapsible-fifth-title]  
local
```

Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|interface  
|link:#interface[interface]
```

```
a|  
  
|ip  
|link:#ip[ip]  
a|IP information to create a new interface.
```

```
|port  
|link:#port[port]  
a|
```

```
|===
```

```
[#peer]  
[.api-collapsible-fifth-title]  
peer
```

Information describing the router to peer with

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|address  
|string  
a|Peer router address
```

```
|asn  
|integer  
a|Autonomous system number of peer
```

```
|===
```

```
[#bgp_peer_group]  
[.api-collapsible-fifth-title]  
bgp_peer_group
```

A BGP peer group between a local network interface and a router, for the purpose of announcing VIP interface locations for SVMs in this IPspace.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name is supplied on input.

|local
|link:#local[local]
a|Information describing the local interface that is being used to peer
with a router using BGP. On a POST operation, an existing BGP interface is
used by specifying the interface, or create a new one by specifying the
port and IP address.

|name
|string
a|Name of the peer group

|peer
|link:#peer[peer]
a|Information describing the router to peer with

|state
|string
a|State of the peer group

|uuid
|string
a|UUID of the peer group

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```

|===

//end collapsible .Definitions block
====

[[ID4cda8c3d94a80c8ea19ea90ab24c11be]]
= Delete a BGP peer group for VIP

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/network/ip/bgp/peer-groups/{uuid}`#

*Introduced In:* 9.7

Deletes a BGP peer group for VIP.

== Related ONTAP commands

* `network bgp peer-group delete`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|UUID of the peer group

|===

== Response

```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description
| 53282019
| Internal error. Failed to remove BGP peer group on node. Wait a few
minutes and try the command again.
|===
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#error[error]
a|
```

```
|===
```

```
.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
```



```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]

```

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[IDbae9a6257bd3b5d41e97446ec1162a05]]

= Retrieve details of a BGP peer group for VIP

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/network/ip/bgp/peer-groups/{uuid}`#

Introduced In: 9.7

Retrieves details of a BGP peer group for VIP.

== Related ONTAP commands

* `network bgp peer-group show`

== Parameters

[cols=5*,options=header]

|===

|Name

```

|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|UUID of the peer group

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name is supplied on input.

|local
|link:#local[local]
a|Information describing the local interface that is being used to peer
with a router using BGP. On a POST operation, an existing BGP interface is
used by specifying the interface, or create a new one by specifying the
port and IP address.

|name
|string
a|Name of the peer group

```

```
|peer
|link:#peer[peer]
a|Information describing the router to peer with
```

```
|state
|string
a|State of the peer group
```

```
|uuid
|string
a|UUID of the peer group
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "local": {
    "interface": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  }
}
```

```

    },
    "port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "elb",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "bgpv4peer",
  "peer": {
    "address": "10.10.10.7"
  },
  "state": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{

```

```

"error": {
  "arguments": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

|===

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type

```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#ipspace]
```

```
[.api-collapsible-fifth-title]
```

```
ipspace
```

Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|IPspace name
```

```
|uuid
```

```
|string
```

```
a|IPspace UUID
```

```
|===
```

```
[#ip]
```

```
[.api-collapsible-fifth-title]
```

```
ip
```

IP information

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|===

[#interface]
[.api-collapsible-fifth-title]
interface

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ip
|link:#ip[ip]
a|IP information

|name
|string
a|The name of the interface.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|===

[#ip]
[.api-collapsible-fifth-title]

```


ip

IP information to create a new interface.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|address

|string

a|IPv4 or IPv6 address

|netmask

|string

a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

|===

[#node]

[.api-collapsible-fifth-title]

node

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Name of node on which the port is located.

|===

[#port]

[.api-collapsible-fifth-title]

port

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#local]
```

```
[.api-collapsible-fifth-title]
```

```
local
```

Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|interface
```

```
|link:#interface[interface]
```

```
a|
```

```
|port
```

```
|link:#port[port]
```

```
a|
```

```
|===
```

```
[#peer]
```

```
[.api-collapsible-fifth-title]
```

```
peer
```

Information describing the router to peer with

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|address
```

```
|string
```

```
a|Peer router address
```

```
|asn
```

```
|integer
```

```
a|Autonomous system number of peer
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

```
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDe7bcb9e3fcc18359efad6bee73897a19]]

```

= Update a BGP peer group for VIP

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-  
block]#`/network/ip/bgp/peer-groups/{uuid}`#
```

Introduced In: 9.7

Updates a BGP peer group for VIP.

== Related ONTAP commands

```
* `network bgp peer-group modify`  
* `network bgp peer-group rename`
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the peer group
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|local
```

```
|link:#local[local]
```

```
a|Information describing the local interface that is being used to peer
```

with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

```
|name
|string
a|Name of the peer group
```

```
|peer
|link:#peer[peer]
a|Information describing the router to peer with
```

```
|state
|string
a|State of the peer group
```

```
|uuid
|string
a|UUID of the peer group
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "name": "bgpv4peer",
  "peer": {
    "address": "10.10.10.7"
  },
  "state": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====
```

```
== Response
```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

1967171	
---------	--

	Internal error. Fail to access or update BGP peer group. Retry the command, if necessary.
--	-------------------------------------------------------------------------------------------

53281998	
----------	--

	Failed to rename the BGP peer group because that name is already assigned to a different peer group in the IPspace.
--	---------------------------------------------------------------------------------------------------------------------

53282006	
----------	--

	BGP peer group could not be updated to use a peer address because the value provided is not a valid peer address. If necessary, try the command again with a routable host address.
--	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

53282007	
----------	--

	BGP peer group could not be updated to use a peer address because the address represents a different address family to the address of the associated BGP LIF. If necessary, try the command again with a matching address family.
--	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

53282018	
----------	--

	Failed to create BGP peer group because an existing peer group has already established a BGP session between LIF and peer address. If necessary, try the command again with a different BGP LIF or a different peer address.
--	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

=====

[#href]

[.api-collapsible-fifth-title]

href

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#ipspace]
```

```
[.api-collapsible-fifth-title]
```

```
ipspace
```

Either the UUID or name is supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|IPspace name
```

```
|uuid
```

```
|string
```

```
a|IPspace UUID
```

```
|===
```

```
[#ip]
```

```
[.api-collapsible-fifth-title]
```

```
ip
```


IP information

```
[#interface]
[.api-collapsible-fifth-title]
interface
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|ip
|link:#ip[ip]
a|IP information
```

```
|name
|string
a|The name of the interface.
```

```
|uuid
|string
a|The UUID that uniquely identifies the interface.
```

```
|===
```

```
[#ip]
[.api-collapsible-fifth-title]
ip
```

IP information to create a new interface.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|address
|string
a|IPv4 or IPv6 address
```

```
|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
```

```
|string
```

```
a|Name of node on which the port is located.
```

```
|===
```

```
[#port]
[.api-collapsible-fifth-title]
port
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#local]
[.api-collapsible-fifth-title]
local
```

Information describing the local interface that is being used to peer with a router using BGP. On a POST operation, an existing BGP interface is used by specifying the interface, or create a new one by specifying the port and IP address.

```
[#peer]
[.api-collapsible-fifth-title]
peer
```

Information describing the router to peer with

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|address
|string
a|Peer router address
```

```
|===
```

```
[#bgp_peer_group]
[.api-collapsible-fifth-title]
bgp_peer_group
```

A BGP peer group between a local network interface and a router, for the purpose of announcing VIP interface locations for SVMs in this IPspace.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|local
|link:#local[local]
a|Information describing the local interface that is being used to peer
with a router using BGP. On a POST operation, an existing BGP interface is
used by specifying the interface, or create a new one by specifying the
port and IP address.

|name
|string
a|Name of the peer group

|peer
|link:#peer[peer]
a|Information describing the router to peer with

|state
|string
a|State of the peer group

|uuid
|string
a|UUID of the peer group

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code

```

```

|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
:leveloffset: -1
```

= Manage network IP interfaces

```
:leveloffset: +1
```

```
[[ID271145f24952492cf373afb01e4692f0]]
```

= Network IP interfaces endpoint overview

== Overview

The following operations are supported:

- * Creation: POST network/ip/interfaces
- * Collection Get: GET network/ip/interfaces
- * Instance Get: GET network/ip/interfaces/{uuid}
- * Instance Patch: PATCH network/ip/interfaces/{uuid}
- * Instance Delete: DELETE network/ip/interfaces/{uuid}

== Retrieving network interface information

The IP interfaces GET API retrieves and displays relevant information pertaining to the interfaces configured in the cluster. The response can contain a list of multiple interfaces or a specific interface. The fields returned in the response vary for different interfaces and configurations.

== Examples

=== Retrieving all interfaces in the cluster

The following example shows the list of all interfaces configured in a cluster.

```
'''
```

```
----
```

```
# The API:  
/api/network/ip/interfaces
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/interfaces" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "14531286-59fc-11e8-ba55-005056b4340f",
      "name": "user-cluster-01_mgmt1",
      "_links": {
        "self": {
          "href": "/api/network/ip/interfaces/14531286-59fc-11e8-ba55-
005056b4340f"
        }
      }
    },
    {
      "uuid": "145318ba-59fc-11e8-ba55-005056b4340f",
      "name": "user-cluster-01_clus2",
      "_links": {
        "self": {
          "href": "/api/network/ip/interfaces/145318ba-59fc-11e8-ba55-
005056b4340f"
        }
      }
    },
    {
      "uuid": "14531e45-59fc-11e8-ba55-005056b4340f",
      "name": "user-cluster-01_clus1",
      "_links": {
        "self": {
          "href": "/api/network/ip/interfaces/14531e45-59fc-11e8-ba55-
005056b4340f"
        }
      }
    },
    {
      "uuid": "245979de-59fc-11e8-ba55-005056b4340f",
      "name": "cluster_mgmt",
      "_links": {
        "self": {
          "href": "/api/network/ip/interfaces/245979de-59fc-11e8-ba55-
005056b4340f"
        }
      }
    }
  ]
}
```

```

    },
    {
      "uuid": "c670707c-5a11-11e8-8fcb-005056b4340f",
      "name": "lif1",
      "_links": {
        "self": {
          "href": "/api/network/ip/interfaces/c670707c-5a11-11e8-8fcb-005056b4340f"
        }
      }
    }
  ],
  "num_records": 5,
  "_links": {
    "self": {
      "href": "/api/network/ip/interfaces"
    }
  }
}
----

'''

```

=== Retrieving a specific Cluster-scoped interface

The following example shows the response when a specific Cluster-scoped interface is requested. The system returns an error when there is no interface with the requested UUID. SVM information is not returned for Cluster-scoped interfaces.

```

'''
----

# The API:
/api/network/ip/interfaces/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/interfaces/245979de-59fc-11e8-ba55-005056b4340f" -H "accept: application/hal+json"

# The response:
{
  "uuid": "245979de-59fc-11e8-ba55-005056b4340f",
  "name": "cluster_mgmt",
  "ip": {

```



```

    "address": "10.63.41.6",
    "netmask": "18",
    "family": "ipv4",
  },
  "enabled": true,
  "state": "up",
  "scope": "cluster",
  "ipospace": {
    "uuid": "114ecfb5-59fc-11e8-ba55-005056b4340f",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/114ecfb5-59fc-11e8-ba55-005056b4340f"
      }
    }
  },
  "services": [
    "management_core",
    "management_autosupport",
    "management_access"
  ],
  "location": {
    "is_home": true,
    "auto_revert": false,
    "failover": "broadcast_domain_only",
    "node": {
      "uuid": "c1db2904-1396-11e9-bb7d-005056acfcbb",
      "name": "user-cluster-01-a",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/c1db2904-1396-11e9-bb7d-005056acfcbb"
        }
      }
    }
  },
  "port": {
    "uuid": "c84d5337-1397-11e9-87c2-005056acfcbb",
    "name": "e0d",
    "node": {
      "name": "user-cluster-01-a"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/c84d5337-1397-11e9-87c2-005056acfcbb"
      }
    }
  }
}

```

```

    }
  },
  "home_node": {
    "uuid": "c1db2904-1396-11e9-bb7d-005056acfcbb",
    "name": "user-cluster-01-a",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/c1db2904-1396-11e9-bb7d-005056acfcbb"
      }
    }
  },
  "home_port": {
    "uuid": "c84d5337-1397-11e9-87c2-005056acfcbb",
    "name": "e0d",
    "node": {
      "name": "user-cluster-01-a"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/c84d5337-1397-11e9-87c2-005056acfcbb"
      }
    }
  },
  "service_policy": {
    "uuid": "9e0f4151-141b-11e9-851e-005056ac1ce0",
    "name": "default-management"
  },
  "vip": false,
  "_links": {
    "self": {
      "href": "/api/network/ip/interfaces/245979de-59fc-11e8-ba55-005056b4340f"
    }
  }
}
----

'''

```

=== Retrieving a specific SVM-scoped interface using a filter

The following example shows the response when a specific SVM-scoped interface is requested. The SVM object is only included for SVM-scoped interfaces.

```

'''

----

# The API:
/api/network/ip/interfaces

# The call:
curl -X GET "https://<mgmt-
ip>/api/network/ip/interfaces?name=lif1&fields=*" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "c670707c-5a11-11e8-8fcb-005056b4340f",
      "name": "lif1",
      "ip": {
        "address": "10.10.10.11",
        "netmask": "24",
        "family": "ipv4",
      },
      "enabled": true,
      "state": "up",
      "scope": "svm",
      "ipspace": {
        "uuid": "114ecfb5-59fc-11e8-ba55-005056b4340f",
        "name": "Default",
        "_links": {
          "self": {
            "href": "/api/network/ipspaces/114ecfb5-59fc-11e8-ba55-
005056b4340f"
          }
        }
      },
    },
    {
      "svm": {
        "uuid": "c2134665-5a11-11e8-8fcb-005056b4340f",
        "name": "user_vs0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/c2134665-5a11-11e8-8fcb-005056b4340f"
          }
        }
      }
    }
  ],
}

```

```

"services": [
  "data_core",
  "data_nfs",
  "data_cifs",
  "data_flexcache"
],
"location": {
  "is_home": true,
  "auto_revert": false,
  "failover": "broadcast_domain_only",
  "node": {
    "uuid": "c1db2904-1396-11e9-bb7d-005056acfcbb",
    "name": "user-cluster-01-a",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/c1db2904-1396-11e9-bb7d-
005056acfcbb"
      }
    }
  },
  "port": {
    "uuid": "c84d5337-1397-11e9-87c2-005056acfcbb",
    "name": "e0d",
    "node": {
      "name": "user-cluster-01-a"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/c84d5337-1397-11e9-87c2-
005056acfcbb"
      }
    }
  },
  "home_node": {
    "uuid": "c1db2904-1396-11e9-bb7d-005056acfcbb",
    "name": "user-cluster-01-a",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/c1db2904-1396-11e9-bb7d-
005056acfcbb"
      }
    }
  },
  "home_port": {
    "uuid": "c84d5337-1397-11e9-87c2-005056acfcbb",
    "name": "e0d",

```

```

    "node": {
      "name": "user-cluster-01-a"
    },
    "_links": {
      "self": {
        "href": "/api/network/ethernet/ports/c84d5337-1397-11e9-87c2-005056acfcbb"
      }
    }
  },
  "service_policy": {
    "uuid": "9e53525f-141b-11e9-851e-005056ac1ce0",
    "name": "default-data-files"
  },
  "vip": false,
  "_links": {
    "self": {
      "href": "/api/network/ip/interfaces/c670707c-5a11-11e8-8fcb-005056b4340f"
    }
  }
},
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/ip/interfaces?name=lif1&fields=*"
  }
}
}
}
-----
'''

```

=== Retrieving specific fields and limiting the output using filters

The following example shows the response when a filter is applied (location.home_port.name=e0a) and only certain fields are requested. Filtered fields are in the output in addition to the default fields and requested fields.

```

'''
-----

```

```

# The API:
/api/network/ip/interfaces

# The call:
curl -X GET "https://<mgmt-
ip>/api/network/ip/interfaces?location.home_port.name=e0a&fields=location.
home_node.name,service_policy.name,ip.address,enabled" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "1d1c9dc8-4f17-11e9-9553-005056ac918a",
      "name": "user-cluster-01-a_clus1",
      "ip": {
        "address": "192.168.170.24"
      },
      "enabled": true,
      "location": {
        "home_node": {
          "name": "user-cluster-01-a"
        },
        "home_port": {
          "name": "e0a"
        }
      },
      "service_policy": {
        "name": "default-cluster"
      },
      "_links": {
        "self": {
          "href": "/api/network/ip/interfaces/1d1c9dc8-4f17-11e9-9553-
005056ac918a"
        }
      }
    },
    {
      "uuid": "d07782c1-4f16-11e9-86e7-005056ace7ee",
      "name": "user-cluster-01-b_clus1",
      "ip": {
        "address": "192.168.170.22"
      },
      "enabled": true,
      "location": {
        "home_node": {

```

```

        "name": "user-cluster-01-b"
    },
    "home_port": {
        "name": "e0a"
    }
},
"service_policy": {
    "name": "default-cluster"
},
"_links": {
    "self": {
        "href": "/api/network/ip/interfaces/d07782c1-4f16-11e9-86e7-005056ace7ee"
    }
}
],
"num_records": 2,
"_links": {
    "self": {
        "href":
"/api/network/ip/interfaces?location.home_port.name=e0a&fields=location.home_node.name,service_policy.name,ip.address,enabled"
    }
}
}
}
-----
'''

```

== Creating IP interfaces

You can use the IP interfaces POST API to create IP interfaces as shown in the following examples.

'''

== Examples

=== Creating a Cluster-scoped IP interface using names

The following example shows the record returned after the creation of an IP interface on "e0d".

'''

```

----

# The API:
/api/network/ip/interfaces

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ip/interfaces?return_records=true" -H "accept:
application/hal+json" -d '{ "name": "cluster_mgmt", "ip": { "address":
"10.63.41.6", "netmask": "18" }, "enabled": true, "scope": "cluster",
"ipspace": { "name": "Default" }, "location": { "auto_revert": false,
"failover": "broadcast_domain_only", "home_port": { "name": "e0d", "node":
{ "name": "user-cluster-01-a" } } }, "service_policy": { "name": "default-
management" } }'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "245979de-59fc-11e8-ba55-005056b4340f",
      "name": "cluster_mgmt",
      "ip": {
        "address": "10.63.41.6",
        "netmask": "18"
      },
      "enabled": true,
      "scope": "cluster",
      "ipspace": {
        "name": "Default"
      },
      "location": {
        "auto_revert": false,
        "failover": "broadcast_domain_only",
        "home_port": {
          "name": "e0d",
          "node": {
            "name": "user-cluster-01-a"
          }
        },
      },
      "service_policy": {
        "name": "default-management"
      },
      "_links": {
        "self": {

```



```

        "href": "/api/network/ip/interfaces/245979de-59fc-11e8-ba55-005056b4340f"
    }
}
]
}
----
'''

```

=== Creating a SVM-scoped IP interface using a mix of parameter types

The following example shows the record returned after the creation of a IP interface by specifying a broadcast domain as the location.

```

'''
----

# The API:
/api/network/ip/interfaces

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ip/interfaces?return_records=true" -H "accept:
application/hal+json" -d '{ "name": "Data1", "ip": { "address":
"10.234.101.116", "netmask": "255.255.240.0" }, "enabled": true, "scope":
"svm", "svm": { "uuid": "137f3618-1e89-11e9-803e-005056a7646a" },
"location": { "auto_revert": true, "broadcast_domain": { "name": "Default"
} }, "service_policy": { "name": "default-data-files" } }'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "80d271c9-1f43-11e9-803e-005056a7646a",
      "name": "Data1",
      "ip": {
        "address": "10.234.101.116",
        "netmask": "20"
      },
      "enabled": true,
      "scope": "svm",
      "svm": {

```

```

    "uuid": "137f3618-1e89-11e9-803e-005056a7646a",
    "name": "vs0",
    "_links": {
      "self": {
        "href": "/api/svm/svms/137f3618-1e89-11e9-803e-005056a7646a"
      }
    }
  },
  "location": {
    "auto_revert": true
  },
  "service_policy": {
    "name": "default-data-files"
  },
  "_links": {
    "self": {
      "href": "/api/network/ip/interfaces/80d271c9-1f43-11e9-803e-005056a7646a"
    }
  }
}
]
}
-----
'''

```

=== Creating a Cluster-scoped IP interface without specifying the scope parameter

The following example shows the record returned after creating an IP interface on "e0d" without specifying the scope parameter. The scope is "cluster" if an "svm" is not specified.

```

'''
-----

# The API:
/api/network/ip/interfaces

# The call:
curl -X POST "https://<mgmt-
ip>/api/network/ip/interfaces?return_records=true" -H "accept:
application/hal+json" -d '{ "name": "cluster_mgmt", "ip": { "address":
"10.63.41.6", "netmask": "18" }, "enabled": true, "ipspace": { "name":

```

```
"Default" }, "location": { "auto_revert": false, "home_port": { "name":
"e0d", "node": { "name": "user-cluster-01-a" } } }, "service_policy": {
"name": "default-management" } }'
```

The response:

```
{
"num_records": 1,
"records": [
  {
    "uuid": "245979de-59fc-11e8-ba55-005056b4340f",
    "name": "cluster_mgmt",
    "ip": {
      "address": "10.63.41.6",
      "netmask": "18"
    },
    "enabled": true,
    "scope": "cluster",
    "ipspace": {
      "name": "Default"
    },
    "location": {
      "auto_revert": false,
      "home_port": {
        "name": "e0d",
        "node": {
          "name": "user-cluster-01-a"
        }
      }
    },
    "service_policy": {
      "name": "default-management"
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/interfaces/245979de-59fc-11e8-ba55-
005056b4340f"
      }
    }
  }
]
}
-----

'''
```

=== Creating an SVM-scoped IP interface without specifying the scope

parameter

The following example shows the record returned after creating an IP interface on "e0d" without specifying the scope parameter. The scope is "svm" if the "svm" field is specified.

'''

The API:

/api/network/ip/interfaces

The call:

```
curl -X POST "https://<mgmt-  
ip>/api/network/ip/interfaces?return_records=true" -H "accept:  
application/hal+json" -d '{ "name": "Data1", "ip": { "address":  
"10.234.101.116", "netmask": "255.255.240.0" }, "enabled": true, "svm": {  
"uuid": "137f3618-1e89-11e9-803e-005056a7646a" }, "location": {  
"auto_revert": true, "broadcast_domain": { "name": "Default" } },  
"service_policy": { "name": "default-data-files" } }'
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "uuid": "80d271c9-1f43-11e9-803e-005056a7646a",  
      "name": "Data1",  
      "ip": {  
        "address": "10.234.101.116",  
        "netmask": "20"  
      },  
      "enabled": true,  
      "scope": "svm",  
      "svm": {  
        "uuid": "137f3618-1e89-11e9-803e-005056a7646a",  
        "name": "vs0",  
        "_links": {  
          "self": {  
            "href": "/api/svms/137f3618-1e89-11e9-803e-005056a7646a"  
          }  
        }  
      },  
      "location": {  
        "auto_revert": true
```

```

    },
    "service_policy": {
        "name": "default-data-files"
    },
    "_links": {
        "self": {
            "href": "/api/network/ip/interfaces/80d271c9-1f43-11e9-803e-005056a7646a"
        }
    }
}
]
}
-----
'''

```

== Updating IP interfaces

You can use the IP interfaces PATCH API to update the attributes of an IP interface.

'''

== Examples

=== Updating the auto revert flag of an IP interface

The following example shows how the PATCH request changes the auto revert flag to 'false'.

'''

The API:

/api/network/ip/interfaces/{uuid}

The call:

```

curl -X PATCH "https://<mgmt-ip>/api/network/ip/interfaces/80d271c9-1f43-11e9-803e-005056a7646a" -H "accept: application/hal+json" -d '{
"location": { "auto_revert": "false" } }'
{
}

```

'''

=== Updating the service policy of an IP interface

The following example shows how the PATCH request changes the service policy to 'default-management'.

'''

The API:

/api/network/ip/interfaces/{uuid}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/network/ip/interfaces/80d271c9-1f43-11e9-803e-005056a7646a" -H "accept: application/hal+json" -d '{
"service_policy": { "name": "default-management" } }'
```

'''

== Deleting IP interfaces

You can use the IP interfaces DELETE API to delete an IP interface in the cluster.

'''

== Example

=== Deleting an IP Interface

The following DELETE request deletes a network IP interface.

'''

The API:

/api/network/ip/interfaces/{uuid}

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/network/ip/interfaces/80d271c9-1f43-11e9-803e-005056a7646a"
{
```

```

}
----

'''

[[ID79fe646b8587f5be7cda09ad77764b97]]
= Retrieve all IP interface details

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ip/interfaces`#

*Introduced In:* 9.6

Retrieves the details of all IP interfaces.

== Related ONTAP Commands

* `network interface show`

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|location.node.name
|string
|query
|False
a|Filter by location.node.name

|location.node.uuid
|string
|query
|False
a|Filter by location.node.uuid

```

```
|location.auto_revert
|boolean
|query
|False
a|Filter by location.auto_revert
```

```
|location.is_home
|boolean
|query
|False
a|Filter by location.is_home
```

```
|location.failover
|string
|query
|False
a|Filter by location.failover
```

```
|location.home_node.name
|string
|query
|False
a|Filter by location.home_node.name
```

```
|location.home_node.uuid
|string
|query
|False
a|Filter by location.home_node.uuid
```

```
|location.home_port.name
|string
|query
|False
a|Filter by location.home_port.name
```

```
|location.home_port.node.name
|string
|query
```



```
|False  
a|Filter by location.home_port.node.name
```

```
|location.home_port.uuid  
|string  
|query  
|False  
a|Filter by location.home_port.uuid
```

```
|location.port.name  
|string  
|query  
|False  
a|Filter by location.port.name
```

```
|location.port.node.name  
|string  
|query  
|False  
a|Filter by location.port.node.name
```

```
|location.port.uuid  
|string  
|query  
|False  
a|Filter by location.port.uuid
```

```
|enabled  
|boolean  
|query  
|False  
a|Filter by enabled
```

```
|vip  
|boolean  
|query  
|False  
a|Filter by vip
```

```
|scope
```

```
|string
|query
|False
a|Filter by scope
```

```
|services
|string
|query
|False
a|Filter by services
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|ipSPACE.name
|string
|query
|False
a|Filter by ipSPACE.name
```

```
|ipSPACE.uuid
|string
|query
|False
a|Filter by ipSPACE.uuid
```

```
|metric.status
|string
|query
|False
a|Filter by metric.status
```

* Introduced in: 9.8

```
|metric.throughput.write
|integer
|query
|False
```

a|Filter by metric.throughput.write

* Introduced in: 9.8

|metric.throughput.read

|integer

|query

|False

a|Filter by metric.throughput.read

* Introduced in: 9.8

|metric.throughput.total

|integer

|query

|False

a|Filter by metric.throughput.total

* Introduced in: 9.8

|metric.timestamp

|string

|query

|False

a|Filter by metric.timestamp

* Introduced in: 9.8

|metric.duration

|string

|query

|False

a|Filter by metric.duration

* Introduced in: 9.8

|uuid

|string

|query

|False

a|Filter by uuid

```
|state
|string
|query
|False
a|Filter by state
```

```
|ip.address
|string
|query
|False
a|Filter by ip.address
```

```
|ip.family
|string
|query
|False
a|Filter by ip.family
```

```
|ip.netmask
|string
|query
|False
a|Filter by ip.netmask
```

```
|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp
```

* Introduced in: 9.8

```
|statistics.status
|string
|query
|False
a|Filter by statistics.status
```

* Introduced in: 9.8

```
|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write
```

* Introduced in: 9.8

```
|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read
```

* Introduced in: 9.8

```
|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total
```

* Introduced in: 9.8

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|service_policy.uuid
|string
|query
|False
a|Filter by service_policy.uuid
```

```
|service_policy.name
|string
|query
|False
a|Filter by service_policy.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1
* Max value: 120
* Min value: 0

```
|order_by
```

```
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|

|records
|array[link:#ip_interface[ip_interface]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  ,
```

```

"num_records": "1",
"records": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7",
      "family": "string",
      "netmask": "24"
    },
    "ipspace": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "exchange",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "location": {
      "failover": "string",
      "home_node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "home_port": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "elb",
        "node": {
          "name": "node1"
        },
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "node": {

```



```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "port": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "elb",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "name": "dataLif1",
  "scope": "string",
  "service_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default-intercluster",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}

```

```

    },
    "services": [
        "data_nfs"
    ],
    "state": "string",
    "statistics": {
        "status": "ok",
        "throughput_raw": {
            "read": "200",
            "total": "1000",
            "write": "100"
        },
        "timestamp": "2017-01-25T11:20:13Z"
    },
    "svm": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

]

}

====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|address
|string
a|IPv4 or IPv6 address

|family
|string
a|IPv4 or IPv6

|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

Either the UUID or name must be supplied on POST for cluster-scoped
objects.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name

|uuid
|string
a|IPspace UUID

```

```
|===
```

```
[#broadcast_domain]  
[.api-collapsible-fifth-title]  
broadcast_domain
```

Broadcast domain UUID along with a readable name.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|name  
|string  
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid  
|string  
a|Broadcast domain UUID
```

```
|===
```

```
[#home_node]  
[.api-collapsible-fifth-title]  
home_node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```

|name
|string
a|

|uuid
|string
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of node on which the port is located.

|===

[#home_port]
[.api-collapsible-fifth-title]
home_port

Port UUID along with readable names. Either the UUID or both names may be
supplied on input.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|name
|string
a|

|node
|link:#node[node]
a|

|uuid
|string
a|

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#port]
[.api-collapsible-fifth-title]
port

```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.


```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, `broadcast_domain` can be specified alone or with `home_node`.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|auto_revert
|boolean
a|
```

```

|failover
|string
a|Defines where an interface may failover.

|home_node
|link:#home_node[home_node]
a|

|home_port
|link:#home_port[home_port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.

|is_home
|boolean
a|

|node
|link:#node[node]
a|

|port
|link:#port[port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the interface.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer

```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

```
metric
```

The most recent sample of I/O metrics for the interface.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|duration
```

```
|string
```

```
a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
```

```
|status
```

```
|string
```

```
a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data".
```

"inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the interface.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#service_policy]
[.api-collapsible-fifth-title]
service_policy
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

The real time I/O statistics for the interface.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|status
```

```
|string
```

```
a|Errors associated with the sample. For example, if the aggregation of
```

data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the interface. This can be used along with
delta time to calculate the rate of throughput bytes per unit of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#ip_interface]
[.api-collapsible-fifth-title]
ip_interface
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|enabled
|boolean
a|The administrative state of the interface.
```

```
|ip
|link:#ip_info[ip_info]
a|IP information
```

```
|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name must be supplied on POST for cluster-scoped objects.
```

```
|location
|link:#location[location]
a|Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node.
```

```

|metric
|link:#metric[metric]
a|The most recent sample of I/O metrics for the interface.

|name
|string
a|Interface name

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".

|service_policy
|link:#service_policy[service_policy]
a|

|services
|array[string]
a|The services associated with the interface.

|state
|string
a|The operational state of the interface.

|statistics
|link:#statistics[statistics]
a|The real time I/O statistics for the interface.

|svm
|link:#svm[svm]
a|Applies only to SVM-scoped objects. Either the UUID or name must be
supplied on POST.

|uuid
|string
a|The UUID that uniquely identifies the interface.

```



```

|vip
|boolean
a|True for a VIP interface, whose location is announced via BGP.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code

```

```

|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDe3f94ada58b33e303eb0086bf802c55d]]
= Create a new cluster-scoped or SVM-scoped interface

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/network/ip/interfaces`#

*Introduced In:* 9.6

Creates a new Cluster-scoped or SVM-scoped interface.

== Required properties

* `name` - Name of the interface to create.
* `ip.address` - IP address for the interface.
* `ip.netmask` - IP subnet of the interface.
* `ipspace.name` or `ipspace.uuid`

*** Required for Cluster-scoped interfaces.

*** Optional for SVM-scoped interfaces.

* `svm.name` or `svm.uuid`

*** Required for an SVM-scoped interface.

```

```

*** Invalid for a Cluster-scoped interface.

* `location.home_port` or `location.home_node` or
`location.broadcast_domain` - One of these properties must be set to a
value to define where the interface will be located.

== Recommended property values

* `service_policy`

*** `for SVM scoped interfaces`

***** `_default-data-files_` for interfaces carrying file-oriented NAS data
traffic

***** `_default-data-blocks_` for interfaces carrying block-oriented SAN
data traffic

***** `_default-management_` for interfaces carrying SVM management requests

*** `for Cluster scoped interfaces`

***** `_default-intercluster_` for interfaces carrying cluster peering
traffic

***** `_default-management_` for interfaces carrying system management
requests

***** `_default-route-announce_` for interfaces carrying BGP peer
connections

== Default property values

If not specified in POST, the following default property values are
assigned:

* `scope`

*** `_svm_` if svm parameter is specified.

*** `_cluster_` if svm parameter is not specified

* `enabled` - `_true_`
* `location.auto_revert` - `_true_`
* `service_policy`

*** `_default-data-files_` if scope is `svm`

```

```

*** _default-management_ if scope is `cluster` and IPspace is not
`Cluster`

*** _default-cluster_ if scope is `svm` and IPspace is `Cluster`

* `failover` - Selects the least restrictive failover policy supported by
all the services in the service policy.

== Related ONTAP commands

* `network interface create`

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean

```

a|The administrative state of the interface.

|ip

|link:#ip_info[ip_info]

a|IP information

|ipspace

|link:#ipspace[ipspace]

a|Either the UUID or name must be supplied on POST for cluster-scoped objects.

|location

|link:#location[location]

a|Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node.

|name

|string

a|Interface name

|scope

|string

a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".

|service_policy

|link:#service_policy[service_policy]

a|

|services

|array[string]

a|The services associated with the interface.

|state

|string

a|The operational state of the interface.

```
|svm
|link:#svm[svm]
a|Applies only to SVM-scoped objects. Either the UUID or name must be
supplied on POST.
```

```
|uuid
|string
a|The UUID that uniquely identifies the interface.
```

```
|vip
|boolean
a|True for a VIP interface, whose location is announced via BGP.
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "ip": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "ipspace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "location": {
    "broadcast_domain": {
      "name": "bd1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "failover": "string",
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "name": "elb",
      "node": {
        "name": "node1"
      }
    }
  }
}
```

```

    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"name": "dataLif1",
"scope": "string",
"service_policy": {
  "name": "default-intercluster",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"services": [
  "data_nfs"
],
"state": "string",
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response

```

Status: 201, Created

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
|
| 1376656
| Cluster interfaces must be in the same subnet. Verify the address and
| netmask are set to the correct values.
|
| 1376963
| Duplicate IP address.
|
| 1966138
| The same IP address may not be used for both a mgmt interface and a
| gateway address.

```

| 1966140
| An interface with the same name already exists.

| 1966987
| The Vserver Broadcast-Domain Home-Node and Home-Port combination is not valid.

| 1967081
| The specified SVM must exist in the specified IPspace.

| 1967082
| The specified ipspace.name does not match the IPspace name of ipspace.uuid.

| 1967102
| POST operation might have left configuration in an inconsistent state. Check the configuration.

| 1967106
| The specified location.home_port.name does not match the specified port name of location.home_port.uuid.

| 1967107
| The location.home_port.uuid specified is not valid.

| 1967108
| The specified location.home_node.name does not match the node name of location.home_node.uuid.

| 1967109
| The specified location.home_port.node.name does not match the node name of location.home_node.uuid.

| 1967110
| The specified location.home_port.node.name does not match location.home_node.name.

| 1967111
| Home node must be specified by at least one location.home_node, location.home_port, or location.broadcast_domain field.

| 1967112
| The specified location.home_node.name does not match the node name of location.home_port.uuid.

| 1967120

| The specified service_policy.name does not match the specified service policy name of service_policy.uuid.

| 1967121

| Invalid service_policy.uuid specified.

| 1967122

| The specified location.broadcast_domain.name does not match the specified broadcast domain name of location.broadcast_domain.uuid.

| 1967123

| The specified IPspace does not match the IPspace name of location.broadcast_domain.uuid.

| 1967124

| The location.broadcast_domain.uuid specified is not valid.

| 1967127

| svm.uuid or svm.name must be provided if scope is "svm".

| 1967128

| ipspace.uuid or ipspace.name must be provided if scope is "cluster".

| 1967129

| The specified location.home_port.uuid is not valid.

| 1967130

| The specified location.home_port.name is not valid.

| 1967131

| The specified location.home_port.uuid and location.home_port.name are not valid.

| 1967135

| The specified location.broadcast_domain.uuid is not valid.

| 1967136

| The specified location.broadcast_domain.name (and ipspace name) is not valid.

| 1967137

| The specified location.broadcast_domain.uuid and location.broadcast_domain.name (and IPspace name) are not valid.

| 1967145

| The specified location.failover is not valid.

```
| 1967146
| The specified svm.name is not valid.

| 1967147
| The specified svm.uuid is not valid.

| 1967153
| No suitable port exists on location.home_node to host the interface.

| 1967154
| Interfaces cannot be created on ports that are down. If a broadcast
domain is specified, ensure that it contains at least one port that is
operationally up.

| 1967381
| Post VIP interfaces requires an effective cluster version of 9.7 or
later.

| 1967382
| VIP interfaces only reside in SVM scope.

| 1967383
| Neither location.home_port.uuid or location.home_port.name should be set
with vip=true.

| 1967384
| Failed to create VIP interface because the home node does not have
active BGP sessions to support Virtual IP (VIP) traffic.

| 1967385
| VIP interfaces with an IPv4 address must use ip.netmask=32. VIP
interfaces with an IPv6 address must use ip.netmask=128.

| 1967387
| The specified IP address is in use by a subnet in this IPspace.

| 5373966
| An iSCSI interface cannot be created in an SVM configured for NVMe.

| 53281065
| The service_policy does not exist in the SVM.
|===
```

== Definitions

[.api-def-first-level]

```

.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.

```

```
|===
```

```
[#ipspace]  
[.api-collapsible-fifth-title]  
ipospace
```

Either the UUID or name must be supplied on POST for cluster-scoped objects.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|IPspace name
```

```
|uuid  
|string  
a|IPspace UUID
```

```
|===
```

```
[#broadcast_domain]  
[.api-collapsible-fifth-title]  
broadcast_domain
```

Broadcast domain UUID along with a readable name.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string
```

```
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid
```

```
|string
```

```
a|Broadcast domain UUID
```

```
|===
```

```
[#home_node]
```

```
[.api-collapsible-fifth-title]
```

```
home_node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#node]
```

```
[.api-collapsible-fifth-title]
```

```
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|Name of node on which the port is located.
```

```
|===
```

```
[#home_port]  
[.api-collapsible-fifth-title]  
home_port
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|
```

```
|node  
|link:#node[node]  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#node]  
[.api-collapsible-fifth-title]  
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#port]
```

```
[.api-collapsible-fifth-title]
```

```
port
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#location]
```

```
[.api-collapsible-fifth-title]
```

```
location
```

Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, `broadcast_domain` can be specified alone or with `home_node`.

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|auto_revert
|boolean
a|

|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name.

|failover
|string
a|Defines where an interface may failover.

|home_node
|link:#home_node[home_node]
a|

|home_port
|link:#home_port[home_port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.

|is_home
|boolean
a|

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the interface.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

The most recent sample of I/O metrics for the interface.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|status
|string
```

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two

collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the interface.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#service_policy]
[.api-collapsible-fifth-title]
service_policy
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

The real time I/O statistics for the interface.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"inconsistent_delta_time" is encountered when the time between two
```

collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

```

[#ip_interface]
[.api-collapsible-fifth-title]
ip_interface

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|The administrative state of the interface.

|ip
|link:#ip_info[ip_info]
a|IP information

|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name must be supplied on POST for cluster-scoped
objects.

|location
|link:#location[location]
a|Current or home location can be modified. Specifying a port implies a
node. Specifying a node allows an appropriate port to be automatically
selected. Ports are not valid and not shown for VIP interfaces. For POST,
broadcast_domain can be specified alone or with home_node.

|name
|string
a|Interface name

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".

|service_policy

```

```

|link:#service_policy[service_policy]
a|

|services
|array[string]
a|The services associated with the interface.

|state
|string
a|The operational state of the interface.

|svm
|link:#svm[svm]
a|Applies only to SVM-scoped objects. Either the UUID or name must be
supplied on POST.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|vip
|boolean
a|True for a VIP interface, whose location is announced via BGP.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDc0cac0ce60fb171be60240c2f5f3f19e]]

```

= Delete an IP interface

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/network/ip/interfaces/{uuid}`#

Introduced In: 9.6

Deletes an IP interface.

== Related ONTAP commands

* `network interface delete`

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|IP interface UUID

|==

== Response

Status: 200, Ok


```

[[IDeb464080131eed6d2b7159b8d10b6d71]]
= Retrieve details for an IP interface

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ip/interfaces/{uuid}`#

*Introduced In:* 9.6

Retrieves details for a specific IP interface.

== Related ONTAP commands

* `network interface show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|IP interface UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|The administrative state of the interface.

|ip
|link:#ip_info[ip_info]
a|IP information

|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name must be supplied on POST for cluster-scoped
objects.

|location
|link:#location[location]
a|Current or home location can be modified. Specifying a port implies a
node. Specifying a node allows an appropriate port to be automatically
selected. Ports are not valid and not shown for VIP interfaces. For POST,
broadcast_domain can be specified alone or with home_node.

|metric
|link:#metric[metric]
a|The most recent sample of I/O metrics for the interface.

|name
|string
a|Interface name

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to

```

"cluster".

|service_policy
|link:#service_policy[service_policy]
a|

|services
|array[string]
a|The services associated with the interface.

|state
|string
a|The operational state of the interface.

|statistics
|link:#statistics[statistics]
a|The real time I/O statistics for the interface.

|svm
|link:#svm[svm]
a|Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|vip
|boolean
a|True for a VIP interface, whose location is announced via BGP.

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
 "_links": {

```

    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ip": {
    "address": "10.10.10.7",
    "family": "string",
    "netmask": "24"
  },
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "location": {
    "failover": "string",
    "home_node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "elb",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  }
}

```

```

    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "port": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "elb",
    "node": {
      "name": "node1"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"name": "dataLif1",
"scope": "string",
"service_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default-intercluster",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"services": [
  "data_nfs"
],

```

```

"state": "string",
"statistics": {
  "status": "ok",
  "throughput_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{

```

```

"error": {
  "arguments": [
    {
      "code": "string",
      "message": "string"
    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

|===

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|self
|link:#href[href]
a|

|===

[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information


[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address


|family
|string
a|IPv4 or IPv6


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

```

Either the UUID or name must be supplied on POST for cluster-scoped

objects.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|IPspace name
```

```
|uuid
```

```
|string
```

```
a|IPspace UUID
```

```
|===
```

```
[#broadcast_domain]
```

```
[.api-collapsible-fifth-title]
```

```
broadcast_domain
```

Broadcast domain UUID along with a readable name.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|Name of the broadcast domain, scoped to its IPspace
```

```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#home_node]
[.api-collapsible-fifth-title]
home_node
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|Name of node on which the port is located.
```

```
|===
```

```
[#home_port]  
[.api-collapsible-fifth-title]  
home_port
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|name  
|string  
a|
```

```
|node  
|link:#node[node]  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#node]  
[.api-collapsible-fifth-title]  
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#port]
[.api-collapsible-fifth-title]
port
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|auto_revert
|boolean
a|

|failover
|string
a|Defines where an interface may failover.
```

```
|home_node
|link:#home_node[home_node]
a|

|home_port
|link:#home_port[home_port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.
```

```
|is_home
|boolean
a|
```

```
|node
|link:#node[node]
a|
```

```
|port
|link:#port[port]
```

a|Port UUID along with readable names. Either the UUID or both names may be supplied on input.

|===

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the interface.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

|===

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

The most recent sample of I/O metrics for the interface.

```
[cols=3*,options=header]
```

|===

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"inconsistent_old_data" is returned when one or more nodes do not have the
latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the interface.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#service_policy]
[.api-collapsible-fifth-title]
service_policy

```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```



```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]  
[.api-collapsible-fifth-title]  
statistics
```

The real time I/O statistics for the interface.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|status  
|string
```

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw  
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp  
|string  
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDd98fa3eb8c5dee34ef1f7f25b0ec9d4a]]

```

= Update an IP interface

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-  
block]#`/network/ip/interfaces/{uuid}`#
```

Introduced In: 9.6

Updates an IP interface.

== Related ONTAP commands

- * `network interface migrate`
- * `network interface modify`
- * `network interface rename`
- * `network interface revert`

== Parameters

```
[cols=5*,options=header]  
|===
```

Name
Type
In
Required
Description

uuid
string
path
True
a IP interface UUID

```
|===
```

== Request Body

```
[cols=3*,options=header]  
|===
```

Name
Type
Description

enabled

```

|boolean
a|The administrative state of the interface.

|ip
|link:#ip_info[ip_info]
a|IP information

|location
|link:#location[location]
a|Current or home location can be modified. Specifying a port implies a
node. Specifying a node allows an appropriate port to be automatically
selected. Ports are not valid and not shown for VIP interfaces. For POST,
broadcast_domain can be specified alone or with home_node.

|name
|string
a|Interface name

|service_policy
|link:#service_policy[service_policy]
a|

|services
|array[string]
a|The services associated with the interface.

|state
|string
a|The operational state of the interface.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|===

.Example request
[%collapsible%closed]
=====

```

```
[source,json,subs=+macros]
{
  "ip": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "location": {
    "failover": "string",
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_port": {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "port": {
      "name": "e1b",
      "node": {
        "name": "node1"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "dataLif1",
  "service_policy": {
    "name": "default-intercluster",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "services": [
    "data_nfs"
  ],
  "state": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "vip": null
}
====
```

== Response

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 1376963
```

```
| Duplicate IP address.
```

```
| 1966138
```

```
| The same IP address may not be used for both a mgmt interface and a  
gateway address.
```

```
| 1967106
```

```
| The specified location.home_port.name does not match the specified port  
name of location.home_port.uuid.
```

```
| 1967107
```

```
| The specified location.home_port.uuid is not valid.
```

```
| 1967111
```

```
| A home node must be specified by at least one location.home_node,  
location.home_port, or location.broadcast_domain field.
```

```
| 1967113
```

```
| The specified location.port.name does not match the port name of  
location.port.uuid.
```

```
| 1967114
```

```
| The specified location.port.uuid is not valid.
```

```
| 1967115
```

```
| The specified location.node.name does not match the node name of  
location.node.uuid.
```

```
| 1967116
```

```
| The specified location.port.node.name does not match the node name of  
location.node.uuid.
```

```
| 1967117
```

| The specified location.port.node.name does not match location.node.name.

| 1967118

| A node must be specified by at least one location.node or location.port field.

| 1967119

| The specified location.node.name does not match the node name of location.port.uuid.

| 1967120

| The specified service_policy.name does not match the specified service policy name of service_policy.uuid.

| 1967121

| The specified service_policy.uuid is not valid.

| 1967125

| You cannot patch the "location.node" or "location.port" fields to migrate interfaces using the iSCSI data protocol. Instead perform the following PATCH operations on the interface: set the "enabled" field to "false"; change one or more "location.home_port" fields to migrate the interface; and then set the "enabled" field to "true".

| 1967129

| The specified location.home_port.uuid is not valid.

| 1967130

| The specified location.home_port.name is not valid.

| 1967131

| The specified location.home_port.uuid and location.home_port.name are not valid.

| 1967132

| The specified location.port.uuid is not valid.

| 1967133

| The specified location.port.name is not valid.

| 1967134

| The specified location.port.uuid and location.port.name are not valid.

| 1967138

| Cannot patch port for a VIP interface. The specified parameter location.port.uuid is not valid.


```
| 1967139
| Cannot patch port for a VIP interface. The specified parameter
location.port.name is not valid.

| 1967140
| Cannot patch port for a VIP interface. The specified parameters
location.port.uuid and location.port.name are not valid.

| 1967141
| Cannot patch home_port for a VIP interface. The specified parameter
location.home_port.uuid is not valid.

| 1967142
| Cannot patch home_port for a VIP interface. The specified parameter
location.home_port.name is not valid.

| 1967143
| Cannot patch home_port for a VIP interface. The specified parameters
location.home_port.uuid and location.home_port.name are not valid.

| 1967145
| The specified location.failover is not valid.

| 1967153
| No suitable port exists on location.home_node to host the interface.

| 1967380
| Cannot patch home_port for a VIP interface. The specified parameter
location.home_port.node.name is not valid. Consider using
location.home_node.name instead.

| 1967386
| Cannot patch port for a VIP interface. The specified parameter
location.port.node.name is not valid. Consider using location.node.name
instead.

| 1967387
| The specified IP address is in use by a subnet in this IPspace.

| 53281065
| The service_policy does not exist in the SVM.
|===
```

== Definitions

[.api-def-first-level]

```

.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.

```

```
|===
```

```
[#ipspace]  
[.api-collapsible-fifth-title]  
ipospace
```

Either the UUID or name must be supplied on POST for cluster-scoped objects.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|IPspace name
```

```
|uuid  
|string  
a|IPspace UUID
```

```
|===
```

```
[#broadcast_domain]  
[.api-collapsible-fifth-title]  
broadcast_domain
```

Broadcast domain UUID along with a readable name.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string
```

a|Name of the broadcast domain, scoped to its IPspace

|uuid

|string

a|Broadcast domain UUID

|===

[#home_node]

[.api-collapsible-fifth-title]

home_node

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|

|uuid

|string

a|

|===

[#node]

[.api-collapsible-fifth-title]

node

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Name of node on which the port is located.

```
|===
```

```
[#home_port]  
[.api-collapsible-fifth-title]  
home_port
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|
```

```
|node  
|link:#node[node]  
a|
```

```
|uuid  
|string  
a|
```

```
|===
```

```
[#node]  
[.api-collapsible-fifth-title]  
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#port]
```

```
[.api-collapsible-fifth-title]
```

```
port
```

Port UUID along with readable names. Either the UUID or both names may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|node
```

```
|link:#node[node]
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#location]
```

```
[.api-collapsible-fifth-title]
```

```
location
```

Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, `broadcast_domain` can be specified alone or with `home_node`.

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|auto_revert
|boolean
a|

|failover
|string
a|Defines where an interface may failover.

|home_node
|link:#home_node[home_node]
a|

|home_port
|link:#home_port[home_port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.

|is_home
|boolean
a|

|node
|link:#node[node]
a|

|port
|link:#port[port]
a|Port UUID along with readable names. Either the UUID or both names may
be supplied on input.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the interface.

```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

```
metric
```

The most recent sample of I/O metrics for the interface.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|duration
```

```
|string
```

```
a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
```

```
|status
```

```
|string
```

```
a|Errors associated with the sample. For example, if the aggregation of
```


data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the interface.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#service_policy]
[.api-collapsible-fifth-title]
service_policy
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the interface. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

The real time I/O statistics for the interface.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|status
```

```
|string
```

```
a|Errors associated with the sample. For example, if the aggregation of
```

data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the interface. This can be used along with
delta time to calculate the rate of throughput bytes per unit of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

Applies only to SVM-scoped objects. Either the UUID or name must be supplied on POST.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
```

a|The unique identifier of the SVM.

|===

[#ip_interface]

[.api-collapsible-fifth-title]

ip_interface

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|The administrative state of the interface.

|ip

|link:#ip_info[ip_info]

a|IP information

|location

|link:#location[location]

a|Current or home location can be modified. Specifying a port implies a node. Specifying a node allows an appropriate port to be automatically selected. Ports are not valid and not shown for VIP interfaces. For POST, broadcast_domain can be specified alone or with home_node.

|name

|string

a|Interface name

|service_policy

|link:#service_policy[service_policy]

a|

|services

|array[string]

a|The services associated with the interface.

```

|state
|string
a|The operational state of the interface.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|===

//end collapsible .Definitions block
====

[[ID72f3eebb7774fea716d78ec2b185d5f9]]
= Retrieve interface historical performance metrics

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ip/interfaces/{uuid}/metrics`#

*Introduced In:* 9.8

Retrieves historical performance metrics for an interface.

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|status
|string
|query
|False
a|Filter by status

```

```
|throughput.write
|integer
|query
|False
a|Filter by throughput.write
```

```
|throughput.read
|integer
|query
|False
a|Filter by throughput.read
```

```
|throughput.total
|integer
|query
|False
a|Filter by throughput.total
```

```
|timestamp
|string
|query
|False
a|Filter by timestamp
```

```
|duration
|string
|query
|False
a|Filter by duration
```

```
|uuid
|string
|path
|True
a|Unique identifier of the interface.
```

```
|interval
|string
|query
|False
```

a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.

The period for each time range is as follows:

- * 1h: Metrics over the most recent hour sampled over 15 seconds.
- * 1d: Metrics over the most recent day sampled over 5 minutes.
- * 1w: Metrics over the most recent week sampled over 30 minutes.
- * 1m: Metrics over the most recent month sampled over 2 hours.
- * 1y: Metrics over the most recent year sampled over a day.
- * Default value: 1
- * enum: ["1h", "1d", "1w", "1m", "1y"]

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

```

|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#records[records]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resource/link"
    }
  }
}

```



```

    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====

```

```

[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description
```

```
|next
|link:href[href]
a|
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the interface.

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|read
|integer
```

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#records]

[.api-collapsible-fifth-title]

records

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|duration

|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|status

|string

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an

expected monotonically increasing value has decreased in value.
"inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the interface.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

```

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage network IP routes

:leveloffset: +1

[[IDb1fb6c217c277f7c60d3e1b40c3f4ca8]]
= Network IP routes endpoint overview


== Overview

```

This endpoint supports the following operations: GET (collection and instance), POST, and DELETE.

'''

== Retrieving network routes

You can use the IP routes GET API to retrieve and display relevant information pertaining to the routes configured in the cluster. The API retrieves the list of all routes configured in the cluster, or a specific route. The fields that are returned in the response will differ with the configuration.

== Examples

=== Retrieving all routes in the cluster

The following output shows the list of all routes configured in a cluster.

'''

The API:

/api/network/ip/routes

The call:

```
curl -X GET "https://<mgmt-ip>/api/network/ip/routes?fields=*" -H "accept: application/hal+json"
```

The response:

```
{
  "records": [
    {
      "uuid": "5fdffb0b-62f8-11e8-853d-005056b4c971",
      "ipspace": {
        "uuid": "84f4beb2-616c-11e8-a4df-005056b4c971",
        "name": "Default",
        "_links": {
          "self": {
            "href": "/api/network/ipspaces/84f4beb2-616c-11e8-a4df-005056b4c971"
          }
        }
      },
      "svm": {
        "uuid": "3243312c-62f8-11e8-853d-005056b4c971",
```

```

    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/3243312c-62f8-11e8-853d-005056b4c971"
      }
    }
  },
  "scope": "svm",
  "destination": {
    "address": "10.4.3.14",
    "netmask": "18",
    "family": "ipv4"
  },
  "gateway": "10.4.3.1",
  "_links": {
    "self": {
      "href": "/api/network/ip/routes/5fdffb0b-62f8-11e8-853d-005056b4c971"
    }
  }
},
{
  "uuid": "84c128d2-62f9-11e8-853d-005056b4c971",
  "ipspace": {
    "uuid": "cc71aadc-62f7-11e8-853d-005056b4c971",
    "name": "ips1",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/cc71aadc-62f7-11e8-853d-005056b4c971"
      }
    }
  },
  "scope": "cluster",
  "destination": {
    "address": "::",
    "netmask": "0",
    "family": "ipv6"
  },
  "gateway": "fd20:8b1e:b255:814e::1",
  "_links": {
    "self": {
      "href": "/api/network/ip/routes/84c128d2-62f9-11e8-853d-005056b4c971"
    }
  }
}

```



```

},
{
  "uuid": "8cc72bcd-616c-11e8-a4df-005056b4c971",
  "ipspace": {
    "uuid": "84f4beb2-616c-11e8-a4df-005056b4c971",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/84f4beb2-616c-11e8-a4df-005056b4c971"
      }
    }
  },
  "scope": "cluster",
  "destination": {
    "address": "0.0.0.0",
    "netmask": "0",
    "family": "ipv4"
  },
  "gateway": "10.224.64.1",
  "_links": {
    "self": {
      "href": "/api/network/ip/routes/8cc72bcd-616c-11e8-a4df-005056b4c971"
    }
  }
},
{
  "uuid": "d63b6eee-62f9-11e8-853d-005056b4c971",
  "ipspace": {
    "uuid": "84f4beb2-616c-11e8-a4df-005056b4c971",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/84f4beb2-616c-11e8-a4df-005056b4c971"
      }
    }
  },
  "svm": {
    "uuid": "3243312c-62f8-11e8-853d-005056b4c971",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/3243312c-62f8-11e8-853d-005056b4c971"
      }
    }
  }
}

```

```

    }
  },
  "scope": "svm",
  "destination": {
    "address": "fd20:8b1e:b255:814e::",
    "netmask": "64",
    "family": "ipv6"
  },
  "gateway": "fd20:8b1e:b255:814e::1",
  "_links": {
    "self": {
      "href": "/api/network/ip/routes/d63b6eee-62f9-11e8-853d-005056b4c971"
    }
  }
}
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/network/ip/routes?fields=*"
  }
}
}
}
----
'''

```

=== Retrieving a specific Cluster-scoped route

The following output shows the returned response when a specific Cluster-scoped route is requested. The system returns an error if there is no route with the requested UUID. SVM information is not returned for Cluster-scoped routes.

```

'''
----

# The API:
/api/network/ip/routes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/routes/84c128d2-62f9-11e8-853d-005056b4c971?fields=*" -H "accept: application/hal+json"

```

```
# The response:
{
  "uuid": "84c128d2-62f9-11e8-853d-005056b4c971",
  "ipspace": {
    "uuid": "cc71aadc-62f7-11e8-853d-005056b4c971",
    "name": "ips1",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/cc71aadc-62f7-11e8-853d-005056b4c971"
      }
    }
  },
  "scope": "cluster",
  "destination": {
    "address": "::",
    "netmask": "0",
    "family": "ipv6"
  },
  "gateway": "fd20:8b1e:b255:814e::1",
  "_links": {
    "self": {
      "href": "/api/network/ip/routes/84c128d2-62f9-11e8-853d-005056b4c971"
    }
  }
}
----

'''
```

=== Retrieving a specific SVM-scoped route

The following output shows the returned response when a specific SVM-scoped route is requested. The system returns an error if there is no route with the requested UUID. The SVM object is only included for SVM-scoped routes.

```
'''

----

# The API:
/api/network/ip/routes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/routes/d63b6eee-62f9-11e8-853d-005056b4c971?fields=*" -H "accept: application/hal+json"
```

```
# The response:
{
  "uuid": "d63b6eee-62f9-11e8-853d-005056b4c971",
  "ipspace": {
    "uuid": "84f4beb2-616c-11e8-a4df-005056b4c971",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/84f4beb2-616c-11e8-a4df-005056b4c971"
      }
    }
  },
  "svm": {
    "uuid": "3243312c-62f8-11e8-853d-005056b4c971",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/3243312c-62f8-11e8-853d-005056b4c971"
      }
    }
  },
  "scope": "svm",
  "destination": {
    "address": "fd20:8b1e:b255:814e::",
    "netmask": "64",
    "family": "ipv6"
  },
  "gateway": "fd20:8b1e:b255:814e::1",
  "_links": {
    "self": {
      "href": "/api/network/ip/routes/d63b6eee-62f9-11e8-853d-005056b4c971"
    }
  }
}
-----

'''
```

== Creating network routes

You can use the POST API to create an SVM-scoped route by specifying the associated SVM, or a Cluster-scoped route by specifying the associated IPspace.

== Examples

=== Creating a Cluster-scoped route

IPspace is required to create a Cluster-scoped route. If the IPspace is not specified, the route will be created in the Default IPspace. The default destination will be set to "0.0.0.0/0" for IPv4 gateway addresses or "::/0" for IPv6 gateway addresses.

'''

The API:

/api/network/ip/routes

The call:

```
curl -X POST "https://<mgmt-ip>/api/network/ip/routes?return_records=true"
-H "accept: application/json" -d '{ "ipspace": { "name":"ips1" },
"gateway": "10.10.10.1"}'
```

The response:

```
{
"num_records": 1,
"records": [
  {
    "uuid": "ae583c9e-9ac7-11e8-8bc9-005056bbd531",
    "ipspace": {
      "name": "ips1"
    },
    "gateway": "10.10.10.1"
  }
]
}
```

'''

=== Creating an SVM-scoped route

To create an SVM-scoped route, the associated SVM can be identified by either its UUID or name.

'''

The API:

/api/network/ip/routes

```
# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/routes?return_records=true"
-H "accept: application/json" -d '{ "svm": { "name":"vs0" }, "gateway":
"10.10.10.1"}'
```

```
# The response:
```

```
{
  "num_records": 1,
  "records": [
    {
      "uuid": "38805a91-9ac9-11e8-8bc9-005056bbd531",
      "svm": {
        "name": "vs0"
      },
      "gateway": "10.10.10.1"
    }
  ]
}
```

```
----
```

```
'''
```

```
== Deleting network routes
```

You can use the DELETE API to delete a specific route identified by its UUID.

```
== Example
```

```
=== Deleting a specific route
```

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ip/routes/{uuid}
```

```
# The call:
```

```
curl -X DELETE "https://<mgmt-ip>/api/network/ip/routes/38805a91-9ac9-
11e8-8bc9-005056bbd531"
```

```
----
```

```
'''
```

```
[[IDeceddd863131ec65885429145731b8329]]
```

= Retrieve IP routes

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/network/ip/routes`#
```

Introduced In: 9.6

Retrieves the collection of IP routes.

== Related ONTAP commands

* `network route show`

== Parameters

```
[cols=5*,options=header]
```

```
|==
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|scope
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by scope
```

```
|ipspace.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by ipspace.name
```

```
|ipspace.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by ipspace.uuid
```

```
|gateway
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by gateway
```

```
|svm.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by svm.uuid
```

```
|svm.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by svm.name
```

```
|uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by uuid
```

```
|destination.address
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by destination.address
```

```
|destination.family
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by destination.family
```

```
|destination.netmask
```

```
|string
```



```
|query
|False
a|Filter by destination.netmask
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1
* Max value: 120
* Min value: 0

```
|order_by
|array[string]
|query
|False
```

a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|error
|link:#error[error]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#network_route[network_route]]
a|

|===
```

```
.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
```

```

    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "num_records": "1",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "destination": {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24"
      },
      "gateway": "10.1.1.1",
      "ipspace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "exchange",
        "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
      },
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    },
  ],

```

```

        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions

```

```

[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|family
|string
a|IPv4 or IPv6

|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always

```

```
netmask length.
```

```
|===
```

```
[#ipspace]  
[.api-collapsible-fifth-title]  
ipspace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|name  
|string  
a|IPspace name
```

```
|uuid  
|string  
a|IPspace UUID
```

```
|===
```

```
[#svm]  
[.api-collapsible-fifth-title]  
svm
```

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```



```

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#network_route]
[.api-collapsible-fifth-title]
network_route

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|destination
|link:#ip_info[ip_info]
a|IP information

|gateway
|string
a|The IP address of the gateway router leading to the destination.

|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.

```

```

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".

|svm
|link:#svm[svm]
a|

|uuid
|string
a|The UUID that uniquely identifies the route.

|===

//end collapsible .Definitions block
=====

[[ID3fd58dd1069f74f2f324f37c9a08b3e9]]
= Create a cluster-scoped or SVM-scoped static route

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/network/ip/routes`#

*Introduced In:* 9.6

Creates a Cluster-scoped or SVM-scoped static route.

== Required properties

* `gateway` - IP address to route packets to.
* SVM-scoped routes

*** `svm.name` or `svm.uuid` - SVM that route is applied to.

* cluster-scoped routes

*** There are no additional required fields for Cluster-scoped routes.

== Default property values

If not specified in POST, the following default property values are

```

assigned:

* `destination` - `_0.0.0.0/0_` for IPv4 or `_::/0_` for IPv6.
* `ipspace.name`

*** `_Default_` for Cluster-scoped routes.

*** Name of the SVM's IPspace for SVM-scoped routes.

== Related ONTAP commands

* ``network route create``

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|==

== Request Body

[cols=3*,options=header]

|==

|Name

|Type

|Description

|destination

|link:#ip_info[ip_info]

a|IP information

|gateway

|string

a|The IP address of the gateway router leading to the destination.

|ipospace

|link:#ipospace[ipospace]

a|Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The UUID that uniquely identifies the route.

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "destination": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "gateway": "10.1.1.1",
  "ipospace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

====

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#network_route[network_route]]
a|

|===


.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "num_records": "1",
  "records": [
    {
      "destination": {
        "address": "10.10.10.7",
        "netmask": "24"
      },
      "gateway": "10.1.1.1",
      "ipspace": {
        "name": "exchange",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
=====

== Error

```

ONTAP Error Response Codes

|===

| Error Code | Description

| 1966345

| Duplicate route exists.

| 1967080

| The destination.address is missing.

| 1967081

| The specified SVM must exist in the specified IPspace.

| 1967082

| The specified ipspace.uuid and ipspace.name refer to different IPspaces.

| 1967146

| The specified svm.name is not valid.

| 2

| The specified svm.uuid is not valid.

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

=====

[#href]

[.api-collapsible-fifth-title]

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

|string

```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.

[cols=3*,options=header]
|===
|Name

```



```
|Type
|Description

|name
|string
a|IPspace name
```

```
|uuid
|string
a|IPspace UUID
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#network_route]
[.api-collapsible-fifth-title]
network_route
```

```
[cols=3*,options=header]
|===
|Name
|Type
```

```

|Description

|destination
|link:#ip_info[ip_info]
a|IP information


|gateway
|string
a|The IP address of the gateway router leading to the destination.


|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.


|svm
|link:#svm[svm]
a|


|uuid
|string
a|The UUID that uniquely identifies the route.


|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string

```

```

a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDc44544a21682b463068e9b8b6ff9088a]]
= Delete an IP route

```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/network/ip/routes/{uuid}`#
```

Introduced In: 9.6

Deletes a specific IP route.

== Related ONTAP commands

* `network route delete`

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|uuid
|string
|path
|True
a|Route UUID
```

```
|===
```

== Response

Status: 200, Ok

== Error

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[IDf423d3bf4b3a0444cb74a1c5a9173dd6]]
= Retrieve details for an IP route

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ip/routes/{uuid}`#

*Introduced In:* 9.6

Retrieves the details of a specific IP route.

== Related ONTAP commands

* `network route show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|Route UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|destination
|link:#ip_info[ip_info]
a|IP information

|gateway
|string
a|The IP address of the gateway router leading to the destination.

|ipspace
|link:#ipspace[ipspace]
a|Applies to both SVM and cluster-scoped objects. Either the UUID or name
may be supplied on input.

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".

|svm
|link:#svm[svm]
a|

|uuid
|string
a|The UUID that uniquely identifies the route.

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]

```



```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "destination": {
    "address": "10.10.10.7",
    "family": "string",
    "netmask": "24"
  },
  "gateway": "10.1.1.1",
  "ipspace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|family
|string
a|IPv4 or IPv6

|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always

```

```
netmask length.
```

```
|===
```

```
[#ipspace]  
[.api-collapsible-fifth-title]  
ipspace
```

Applies to both SVM and cluster-scoped objects. Either the UUID or name may be supplied on input.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|name  
|string  
a|IPspace name
```

```
|uuid  
|string  
a|IPspace UUID
```

```
|===
```

```
[#svm]  
[.api-collapsible-fifth-title]  
svm
```

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

```

|===

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]

```

```

|===

```

```

|Name
|Type
|Description

```

```

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

```

```

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]

```

```

|===

```

```

|Name

```

```

|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage network IP service policies

:leveloffset: +1

[[IDba2a2ea593c8cae6bd17315d4ac836c3]]
= Network IP service-policies endpoint overview

== Overview

Service policies are named groupings that define what services are

```

supported by an IP interface.

The following operations are supported:

*** Creation: POST network/ip/service-policies

*** Collection Get: GET network/ip/service-policies

*** Instance Get: GET network/ip/service-policies/{uuid}

*** Instance Patch: PATCH network/ip/service-policies/{uuid}

*** Instance Delete: DELETE network/ip/service-polices/{uuid}

== Examples

=== Retrieving all service policies in the cluster

The following output shows the collection of all service policies configured in a 2-node cluster. By default (without 'field=*' parameter), only the UUID and name fields are shown for each entry.

'''

The API:

/api/network/ethernet/ip/service-policies

The call:

```
curl -X GET "https://<mgmt-ip>/api/network/ip/service-policies" -H
"accept: application/hal+json"
```

The response:

```
{
  "records": [
    {
      "uuid": "e4e2f193-c1a3-11e8-bb9d-005056bb88c8",
      "name": "net-intercluster",
      "_links": {
        "self": {
          "href": "/api/network/ip/service-policies/e4e2f193-c1a3-11e8-bb9d-005056bb88c8"
        }
      }
    },
    {
      "uuid": "e4e3f6da-c1a3-11e8-bb9d-005056bb88c8",
```

```

    "name": "net-route-announce",
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/e4e3f6da-c1a3-11e8-bb9d-005056bb88c8"
      }
    },
    {
      "uuid": "e5111111-c1a3-11e8-bb9d-005056bb88c8",
      "name": "vserver-route-announce",
      "_links": {
        "self": {
          "href": "/api/network/ip/service-policies/e5111111-c1a3-11e8-bb9d-005056bb88c8"
        }
      }
    },
    {
      "uuid": "e6111111-c1a3-11e8-bb9d-005056bb88c8",
      "name": "data-route-announce",
      "_links": {
        "self": {
          "href": "/api/network/ip/service-policies/e6111111-c1a3-11e8-bb9d-005056bb88c8"
        }
      }
    }
  ],
  "num_records": 4,
  "_links": {
    "self": {
      "href": "/api/network/ip/service-policies/?return_records=true&return_timeout=15"
    }
  }
}
----

'''

```

=== Retrieving a specific service policy (scope=svm)

The following output displays the response when a specific "svm" scoped service policy is requested. Among other parameters, the response contains the svm parameters associated with the service policy. The system returns


```

an error when there is no service policy with the requested UUID.

'''

----

# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X GET "http://<mgmt-ip>/api/network/ip/service-policies/dad323ff-4ce0-11e9-9372-005056bb91a8?fields=*" -H "accept: application/hal+json"

# The response:
{
  "uuid": "dad323ff-4ce0-11e9-9372-005056bb91a8",
  "name": "default-data-files",
  "scope": "svm",
  "svm": {
    "uuid": "d9060680-4ce0-11e9-9372-005056bb91a8",
    "name": "vs0",
    "_links": {
      "self": {
        "href": "/api/svm/svms/d9060680-4ce0-11e9-9372-005056bb91a8"
      }
    }
  },
  "ipspace": {
    "uuid": "45ec2dee-4ce0-11e9-9372-005056bb91a8",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/45ec2dee-4ce0-11e9-9372-005056bb91a8"
      }
    }
  },
  "services": [
    "data_core",
    "data_nfs",
    "data_cifs",
    "data_flexcache"
  ],
  "_links": {
    "self": {
      "href": "/api/network/ip/service-policies/dad323ff-4ce0-11e9-9372-005056bb91a8"
    }
  }
}

```

```

    }
}
}
----

'''

=== Retrieving a specific service policy (scope=svm) when requesting
commonly used fields

The following output displays the response when commonly used fields are
requested for a specific "svm" scoped service policy. Among other
parameters, the response contains the svm parameters associated with the
service policy. The system returns an error when there is no service
policy with the requested UUID.

'''

----

# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/service-policies/e0889ce6-
1e6a-11e9-89d6-005056bbdc04?fields=name,scope,svm.name,ipspace.name" -H
"accept: application/hal+json"

# The response:
{
  "uuid": "e0889ce6-1e6a-11e9-89d6-005056bbdc04",
  "name": "test_policy",
  "scope": "svm",
  "svm": {
    "name": "vs0"
  },
  "ipspace": {
    "name": "Default"
  },
  "_links": {
    "self": {
      "href": "/api/network/ip/service-policies/e0889ce6-1e6a-11e9-89d6-
005056bbdc04"
    }
  }
}

```

'''

=== Retrieving a specific service policy (scope=cluster)

The following output displays the response when a specific cluster-scoped service policy is requested. The SVM object is not included for cluster-scoped service policies. A service policy with a scope of "cluster" is associated with an IPspace. The system returns an error when there is no service policy with the requested UUID.

'''

The API:

/api/network/ip/service-policies/{uuid}

The call:

curl -X GET "https://<mgmt-ip>/api/network/ip/service-policies/4c6b72b9-0f6c-11e9-875d-005056bb21b8?fields=*" -H "accept: application/hal+json"

The response:

```
{
  "uuid": "4c6b72b9-0f6c-11e9-875d-005056bb21b8",
  "name": "net-intercluster",
  "scope": "cluster",
  "ipspace": {
    "uuid": "4051f13e-0f6c-11e9-875d-005056bb21b8",
    "name": "Default",
    "_links": {
      "self": {
        "href": "/api/network/ipspaces/4051f13e-0f6c-11e9-875d-005056bb21b8"
      }
    }
  },
  "services": [
    "intercluster_core"
  ],
  "_links": {
    "self": {
      "href": "/api/network/ip/service-policies/4c6b72b9-0f6c-11e9-875d-005056bb21b8"
    }
  }
}
```

```

}
----

'''

=== Retrieving a specific service policy (scope=cluster) when requesting
commonly used fields

The following output displays the response when commonly used fields are
requested for a specific "cluster" scoped service policy. The SVM object
is not included for cluster-scoped service policies. A service policy with
a scope of "cluster" is associated with an IPspace. The system returns an
error when there is no service policy with the requested UUID.

'''

----

# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/network/ip/service-policies/4c6b72b9-
0f6c-11e9-875d-005056bb21b8?fields=name,scope,ipspace.name" -H "accept:
application/hal+json"

# The response:
{
  "uuid": "4c6b72b9-0f6c-11e9-875d-005056bb21b8",
  "name": "net-intercluster",
  "scope": "cluster",
  "ipspace": {
    "name": "Default"
  },
  "services": [
    "intercluster_core"
  ],
  "_links": {
    "self": {
      "href": "/api/network/ip/service-policies/4c6b72b9-0f6c-11e9-875d-
005056bb21b8"
    }
  }
}
}
}
----

```

```
'''
```

```
== Creating service policies
```

You can use this API to create an SVM-scoped service policy by specifying the associated SVM, or a cluster-scoped service policy by specifying the associated IPspace. If the scope is not specified, it is inferred from the presence of the IPspace or SVM.

Cluster scoped service policies will operate on the IPspace "Default" unless IPspace is explicitly specified.

```
== Examples
```

```
=== Creating a cluster-scoped service policy
```

The following output displays the response when creating a service policy with a scope of "cluster" and an IPspace of "Default".

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ip/service-policies
```

```
# The call:
```

```
curl -X POST "https://<mgmt-ip>/api/network/ip/service-  
policies?return_records=true" -H "accept: application/json" -d '{ "name":  
"new-policy", "scope": "cluster", "ipspace": { "name":"Default" },  
"services": [ "intercluster_core" ] }'
```

```
# The response:
```

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "uuid": "74139267-f1aa-11e9-b5d7-005056a73e2e",  
      "name": "new-policy",  
      "scope": "cluster",  
      "ipspace": {  
        "uuid": "ba556295-e912-11e9-a1c8-005056a7080e",  
        "name": "Default",  
        "_links": {  
          "self": {  
            "href": "/api/network/ipspaces/ba556295-e912-11e9-a1c8-  
005056a7080e"  
          }  
        }  
      }  
    ]  
  }
```

```

    }
  },
  "_links": {
    "self": {
      "href": "/api/network/ip/service-policies/74139267-f1aa-11e9-b5d7-005056a73e2e"
    }
  },
  "services": [
    "intercluster_core"
  ]
}
]
}
}
----
'''

```

=== Creating a cluster-scoped service policy without specifying IPspace

The following output displays the response when creating a service policy with a scope of "cluster" without specifying an IPspace".

```

'''
----

# The API:
/api/network/ip/service-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/service-
policies?return_records=true" -H "accept: application/json" -d '{ "name":
"new-policy", "scope": "cluster", "services": [ "intercluster_core" ] }'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "74139267-f1aa-11e9-b5d7-005056a73e2e",
      "name": "new-policy",
      "scope": "cluster",
      "ipspace": {
        "uuid": "ba556295-e912-11e9-a1c8-005056a7080e",
        "name": "Default",

```

```

    "_links": {
      "self": {
        "href": "/api/network/ipspaces/ba556295-e912-11e9-a1c8-005056a7080e"
      }
    },
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/74139267-f1aa-11e9-b5d7-005056a73e2e"
      }
    },
    "services": [
      "intercluster_core"
    ]
  }
}
}
}

```

'''

=== Creating a cluster-scoped service policy without specifying scope

The following output displays the response when creating a service policy in the "Default" IPspace without specifying the scope".

'''

The API:

/api/network/ip/service-policies

The call:

```

curl -X POST "https://<mgmt-ip>/api/network/ip/service-policies?return_records=true" -H "accept: application/json" -d '{ "name": "new-policy2", "ipspace.name": "Default", "services": [ "intercluster_core" ] }'

```

The response:

```

{
  "num_records": 1,
  "records": [
    {

```

```

    "uuid": "59439267-f1aa-11e9-b5d7-005056a73e2e",
    "name": "new-policy2",
    "scope": "cluster",
    "ipspace": {
      "uuid": "ba556295-e912-11e9-a1c8-005056a7080e",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/ba556295-e912-11e9-a1c8-
005056a7080e"
        }
      }
    },
    "services": [
      "intercluster_core"
    ],
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/74139267-f1aa-11e9-b5d7-
005056a73e2e"
      }
    }
  ]
}
----

'''

```

=== Creating an SVM-scoped service policy

The following output displays the response when creating a service policy with a scope of "svm" in the SVM "vs0".

```

'''
----

# The API:
/api/network/ip/service-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ip/service-
policies?return_records=true" -H "accept: application/json" -d '{ "name":
"new-policy", "scope": "svm", "svm": { "name":"vs0" }, "services": [
"data-nfs", "data-cifs" ] }'

```



```

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "f3901097-f2c4-11e9-b5d7-005056a73e2e",
      "name": "new-policy",
      "scope": "svm",
      "svm": {
        "uuid": "07df9cee-e912-11e9-a13a-005056a73e2e",
        "name": "vs0",
        "_links": {
          "self": {
            "href": "/api/svm/svms/07df9cee-e912-11e9-a13a-005056a73e2e"
          }
        }
      },
      "ipspace": {
        "uuid": "1d3199d2-e906-11e9-a13a-005056a73e2e",
        "name": "Default",
        "_links": {
          "self": {
            "href": "/api/network/ipspaces/1d3199d2-e906-11e9-a13a-005056a73e2e"
          }
        }
      },
      "services": [
        "data_nfs",
        "data_cifs"
      ],
      "_links": {
        "self": {
          "href": "/api/network/ip/service-policies/f3901097-f2c4-11e9-b5d7-005056a73e2e"
        }
      }
    }
  ]
}
'''

=== Creating an SVM-scoped service policy without specifying scope

```

The following output displays the response when creating a service policy with a SVM of "vs0" without specifying the scope.

```
'''
```

```
----
```

```
# The API:
```

```
/api/network/ip/service-policies
```

```
# The call:
```

```
curl -X POST "https://<mgmt-ip>/api/network/ip/service-  
policies?return_records=true" -H "accept: application/json" -d '{ "name":  
"new-policy", "svm": { "name":"vs0" }, "services": [ "data-nfs", "data-  
cifs" ] }'
```

```
# The response:
```

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "uuid": "f3901097-f2c4-11e9-b5d7-005056a73e2e",  
      "name": "new-policy",  
      "scope": "svm",  
      "svm": {  
        "uuid": "07df9cee-e912-11e9-a13a-005056a73e2e",  
        "name": "vs0",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/07df9cee-e912-11e9-a13a-005056a73e2e"  
          }  
        }  
      },  
      "ipspace": {  
        "uuid": "1d3199d2-e906-11e9-a13a-005056a73e2e",  
        "name": "Default",  
        "_links": {  
          "self": {  
            "href": "/api/network/ipspaces/1d3199d2-e906-11e9-a13a-  
005056a73e2e"  
          }  
        }  
      },  
      "services": [  
        "data_nfs",  
        "data_cifs"
```

```

    ],
    "_links": {
      "self": {
        "href": "/api/network/ip/service-policies/f3901097-f2c4-11e9-b5d7-005056a73e2e"
      }
    }
  }
]
}

```

'''

=== Updating the name of a service policy

The following example displays the command used to update the name of a service policy scoped to a specific "svm". The system returns an error when there is no service policy associated with the UUID or the service policy cannot be renamed.

'''

The API:

/api/network/ip/service-policies/{uuid}

The call:

```

curl -X PATCH "https://<mgmt-ip>/api/network/ip/service-policies/734eaf57-d2fe-11e9-9284-005056acaad4" -d "{ \"name\": \"new-name\" }" -H "accept: application/hal+json"

```

'''

=== Updating the services for a service policy

The following example displays the command used to update the services a service policy contains. The system returns an error when there is no service policy associated with the UUID or the services cannot be applied.

'''

```

# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/network/ip/service-policies/734eaf57-
d2fe-11e9-9284-005056acaad4" -d "{ \"services\": [ \"data-nfs\", \"data-
cifs\" ] }" -H "accept: application/hal+json"
----

'''

=== Deleting a service policy

The following output displays the response for deleting a service policy.

'''

----

# The API:
/api/network/ip/service-policies/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ip/service-
policies/757ed726-bdc1-11e9-8a92-005056a7bf25" -H "accept:
application/hal+json"
----

'''

[[ID929d14462ce702ba5416d5f21de5fa0f]]
= Retrieve service policies

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ip/service-policies`#

*Introduced In:* 9.6

Retrieves a collection of service policies.

== Related ONTAP commands

* `network interface service-policy show`

```

```

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|uuid
|string
|query
|False
a|Filter by uuid

|scope
|string
|query
|False
a|Filter by scope

|services
|string
|query
|False
a|Filter by services

```

```
|ipospace.name
|string
|query
|False
a|Filter by ipospace.name
```

```
|ipospace.uuid
|string
|query
|False
a|Filter by ipospace.uuid
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

```
* Default value: 1
```

```
|return_timeout
```

```

|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|

|records
|array[link:#ip_service_policy[ip_service_policy]]
a|

|===

```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ipospace": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "exchange",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "default-intercluster",
      "scope": "string",
      "services": [
        "data_nfs"
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```



```

]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|self
|link:#href[href]
a|

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name

|uuid
|string
a|IPspace UUID

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links

```

```

|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#ip_service_policy]
[.api-collapsible-fifth-title]
ip_service_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ipspace
|link:#ipspace[ipspace]
a|

|name
|string
a|

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".

|services
|array[string]
a|

```

```

|svm
|link:#svm[svm]
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]

```

a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID47117682d69ae6ece4b32f8029f67ea3]]
= Create a service policy for network interfaces

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/network/ip/service-policies`#

Introduced In: 9.8

Creates a service policy for network interfaces.

== Required properties

* `name` - Name of the service policy to create.
* `ipspace.name` or `ipspace.uuid`

*** Required for cluster-scoped service policies.

*** Optional for SVM-scoped service policies.

* `svm.name` or `svm.uuid`

*** Required for SVM-scoped service policies.

*** Not valid for cluster-scoped service policies.

== Default property values

If not specified in POST, the following default property values are assigned:

* `scope`

*** svm if the svm parameter is specified

*** cluster if the svm parameter is not specified

== Parameters

[cols=5*,options=header]
|==

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.

* Default value:

|==

== Request Body

[cols=3*,options=header]
|==
|Name
|Type
|Description

```

|ipospace
|link:#ipospace[ipospace]
a|

|name
|string
a|

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".

|services
|array[string]
a|

|svm
|link:#svm[svm]
a|

|uuid
|string
a|

|===

```

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```

{
  "ipospace": {
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "default-intercluster",
  "scope": "string",
  "services": [
    "data_nfs"
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}

```



```
},
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|====
| Error Code | Description

| 1967146
| Svm.name does not exist.

| 1967147
| Svm.uuid does not exist.

| 53281929
| Service policies cannot combine block and file services.

| 53281931
| Service policy names cannot start with "default-".

| 53281933
| Cluster-scoped services cannot be added to a SVM-scoped service policy.

| 53281934
| SVM-scoped services cannot be added to a Cluster-scoped service policy.

| 53281935
| Scope is set to "svm" and svm.uuid or svm.name have not been specified.

| 53281936
| The SVM is not in the specified IPspace.

| 53281937
| Svm.uuid and svm.name are not valid parameters when scope is cluster.
```

```
| 53281938
| Svm.uuid or svm.name specify a vserver that does not exist.

| 53281939
| One or more of the svm.uuid, svm.name, ipspace.uuid, and ipspace.name
have invalid values.

| 53281940
| SVM or IPspace has not been specified.

| 53281941
| SVM does not exist.

| 53281944
| Ipspace.name does not exist.

| 53281945
| Ipspace.uuid is not an IPspace.

| 53281946
| Service policy already exists.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===
```

```

[#_links]
[.api-collapsible-fifth-title]
_links
[#ipspace]
[.api-collapsible-fifth-title]
ipspace

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|IPspace name

|uuid
|string
a|IPspace UUID

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

```

|===

[#ip_service_policy]
[.api-collapsible-fifth-title]
ip_service_policy

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ipspace
|link:#ipspace[ipspace]
a|

|name
|string
a|

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".

|services
|array[string]
a|

|svm
|link:#svm[svm]
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[ID4eb12149b06bc377c5b93546f46dac75]]
```

```
= Delete a service policy for network interfaces
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/network/ip/service-policies/{uuid}`#
```

```
*Introduced In:* 9.8
```

```
Deletes a service policy for network interfaces.
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the service policy
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
== Error
```

ONTAP Error Response Codes

|===

| Error Code | Description

| 53281927

| Service policies owned by the system cannot be deleted.

| 53281928

| Service policies assigned to LIFs cannot be deleted.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

=====

[source,json,subs=+macros]

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

```

====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

```



```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID18f971537f8e540ffb3ee8442103827a]]
= Retrieve a service policy

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ip/service-policies/{uuid}`#

*Introduced In:* 9.6

Retrieves a specific service policy.

== Related ONTAP commands

* `network interface service-policy show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required

```

```

|Description

|uuid
|string
|path
|True
a|Service policy UUID


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.


* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.


* Default value: 1
* Max value: 120
* Min value: 0

```

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ipspace
|link:#ipspace[ipspace]
a|

|name
|string
a|

|scope
|string
a|Set to "svm" for interfaces owned by an SVM. Otherwise, set to
"cluster".

|services
|array[string]
a|

|svm
|link:#svm[svm]
a|

|uuid
```

```

|string
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ipospace": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "default-intercluster",
  "scope": "string",
  "services": [
    "data_nfs"
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
=====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#ipospace]
[.api-collapsible-fifth-title]
ipospace

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name

|uuid

```

```

|string
a|IPspace UUID

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```



```
[[ID93512eeacel1a34e8aec795d7901dbd5c]]
```

= Update a service policy for network interfaces

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-  
block]#`/network/ip/service-policies/{uuid}`#
```

Introduced In: 9.8

Updates a service policy for network interfaces.

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the service policy
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|services
```

```
|array[string]
```

```

a|

|uuid
|string
a|

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "name": "default-intercluster",
  "services": [
    "data_nfs"
  ],
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
=====

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

```

ONTAP Error Response Codes

|===
| Error Code | Description
| 53281929
| Service policies cannot combine block and file services.
| 53281930
| Service policies maintained by the system cannot be renamed.
| 53281931
| Service policy names cannot start with "default-".
| 53281933

```

| A Cluster-scoped service cannot be added to a SVM-scoped service policy.

| 53281934

| An SVM-scoped service cannot be added to a Cluster-scoped service policy.

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#href]

[.api-collapsible-fifth-title]

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

|string

a|

|===

[#_links]

[.api-collapsible-fifth-title]

_links

[#ipospace]

[.api-collapsible-fifth-title]

ipospace

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|IPspace name

|uuid

|string

a|IPspace UUID

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#ip_service_policy]

[.api-collapsible-fifth-title]

ip_service_policy

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|

```

|services
|array[string]
a|

|uuid
|string
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]

```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage network IPspaces
```

```
:leveloffset: +1
```

```
[[ID54e50a1d7fb05be2fc67053a891137e8]]
```

```
= Network ipspaces endpoint overview
```

```
== Overview
```

An IPspace is an addressing domain within which each IP address is unique. The same address may appear in a different IPspace, but the matching addresses are considered to be distinct. SVMs and broadcast domains, and therefore IP interfaces and Ethernet ports, are associated with a single IPspace. This endpoint supports the following operations: GET (collection and instance), POST, PATCH, and DELETE.

== Retrieving IPspace information

You can use the IPspaces GET API to retrieve all IPspaces configured in the cluster, including built-in and custom IPspaces, and specifically requested IPspaces.

== Examples

=== Retrieving a list of the IPspaces in the cluster

The following example returns the requested list of IPspaces configured in the cluster.

The API:

/api/network/ipspaces

The call:

```
curl -X GET "https://<mgmt-ip>/api/network/ipspaces?fields=*" -H "accept: application/hal+json"
```

The response:

```
{
  "records": [
    {
      "uuid": "dcc7e79c-5acc-11e8-b9de-005056b42b32",
      "name": "Default",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/dcc7e79c-5acc-11e8-b9de-005056b42b32"
        }
      }
    },
    {
      "uuid": "dfd3c1b2-5acc-11e8-b9de-005056b42b32",
      "name": "Cluster",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/dfd3c1b2-5acc-11e8-b9de-005056b42b32"
        }
      }
    }
  ],
}
```

```

{
  "uuid": "dedec1be-5aec-1eee-beee-0eee56be2b3e",
  "name": "IpSPACE1",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/dedec1be-5aec-1eee-beee-0eee56be2b3e"
    }
  }
},
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/network/ipspaces?fields=*"
  }
}
}
----
'''

```

=== Retrieving a specific IPspace in the cluster

The following example returns the specific IPspace requested. The system returns an error if there is no IPspace with the requested UUID.

```

'''
----

```

The API:

/api/network/ipspaces/{uuid}

The call:

```

curl -X GET "https://<mgmt-ip>/api/network/ipspaces/dfd3c1b2-5acc-11e8-b9de-005056b42b32?fields=*" -H "accept: application/hal+json"

```

The response:

```

{
  "uuid": "dcc7e79c-5acc-11e8-b9de-005056b42b32",
  "name": "Default",
  "_links": {
    "self": {
      "href": "/api/network/ipspaces/dcc7e79c-5acc-11e8-b9de-005056b42b32"
    }
  }
}

```



```

    }
}
----

'''

== Creating IPspaces

You can use the network IPspaces POST API to create IPspaces.

'''

== Example

=== Creating an IPspace

The following output displays the record returned after the creation of an
IPspace with the name "ipspace1".

'''

----

# The API:
/api/network/ipspaces

# The call:
curl -X POST "https://<mgmt-ip>/api/network/ipspaces?return_records=true"
-H "accept: application/hal+json" -d '{"name": "ipspace2"}'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "4165655e-0528-11e9-bd68-005056bb046a",
      "name": "ipspace2",
      "_links": {
        "self": {
          "href": "/api/network/ipspaces/4165655e-0528-11e9-bd68-
005056bb046a"
        }
      }
    }
  ]
}
----

```

'''

== Updating IPspaces

You can use the IPspaces PATCH API to update the attributes of the IPspace.

'''

== Example

=== Updating the name of an IPspace

The following PATCH request is used to update the name of the IPspace from "ipspace2" to "ipspace20".

'''

The API:

/api/network/ipspaces/{uuid}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/network/ipspaces/4165655e-0528-11e9-  
bd68-005056bb046a" -H "accept: application/hal+json" -d "{  \"name\":  
\"ipspace20\"}"
```

'''

== Deleting IPspaces

You can use the IPspaces DELETE API to delete an IPspace.

'''

== Example

=== Deleting an IPspace

The following DELETE request is used to delete an IPspace.

'''

```

# The API:
/api/network/ipspaces/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/network/ipspaces/4165655e-0528-11e9-
bd68-005056bb046a" -H "accept: application/hal+json" -H "Content-Type:
application/json"
----

'''

[[ID48e044460f8cfa84e912c45ee2a7cbf1]]
= Retrieve IPspaces for a cluster

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ipspaces`#

*Introduced In:* 9.6

Retrieves a collection of IPspaces for the entire cluster.

== Related ONTAP commands

* `network ipspace show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|name
|string
|query
|False
a|Filter by name

```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1
* Max value: 120
* Min value: 0

```

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|

|records
|array[link:#ipospace[ipospace]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  }
}

```

```

},
"num_records": "1",
"records": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "ipspacel",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ]
  },

```

```

    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next
|link:#href[href]
a|

```

```

|self
|link:#href[href]

```

```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name

|uuid
|string
a|The UUID that uniquely identifies the IPspace.

|===

```



```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[ID63a733a61e090ac237032cc12555743d]]

= Create a new domain with unique IP addresses

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/network/ipspaces`#

Introduced In: 9.6

Creates a new domain within which IP addresses are unique. SVMs, ports, and networks are scoped to a single IPspace.

== Required properties

* `name` - Name of the IPspace to create.

== Related ONTAP commands

* `network ipspace create`

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

```

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|IPspace name

|uuid
|string
a|The UUID that uniquely identifies the IPspace.

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "name": "ipspace1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
=====

== Response

```

Status: 201, Created

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 1967102
```

```
| A POST operation might have left the configuration in an inconsistent  
state. Check the configuration.
```

```
|===
```

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 9240591
```

```
| The name is not valid. The name is already in use by a cluster node,  
Vserver, or it is the name of the local cluster.
```

```
|===
```

== Definitions

```
[.api-def-first-level]
```

```
.See Definitions
```

```
[%collapsible%closed]
```

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

|===

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#ipspace]
[.api-collapsible-fifth-title]
ipspace
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|name
|string
a|IPspace name
```

```
|uuid
|string
a|The UUID that uniquely identifies the IPspace.
```

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
```

```

a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID14e583817e0a4d09607876fd04cc474a]]
= Delete an IPspace object

```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/network/ipspaces/{uuid}`#
```

Introduced In: 9.6

Deletes an IPspace object.

== Related ONTAP commands

* `network ipspace delete`

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|IPspace UUID
```

```
|===
```

== Response

Status: 200, Ok

```

[[ID9f49a8732125cde3d93e5dcc2f09af4]]
= Retrieve information about an IPspace

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/ipspaces/{uuid}`#

*Introduced In:* 9.6

Retrieves information about a specific IPspace.

== Related ONTAP commands

* `network ipspace show`

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|IPspace UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name

|uuid
|string
a|The UUID that uniquely identifies the IPspace.

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "ipspace1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name

```

```

|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID2de75470e248fc90a6427f88dd2e8ed9]]
= Update an IPspace object

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/network/ipspaces/{uuid}`#

*Introduced In:* 9.6

Updates an IPspace object.

```

== Related ONTAP commands

* `network ipspace rename`

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|IPspace UUID

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|IPspace name

|uuid

|string

a|The UUID that uniquely identifies the IPspace.

|===

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "name": "ipspacel",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#ipspacel]
[.api-collapsible-fifth-title]
ipspacel
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|IPspace name


|uuid
|string
a|The UUID that uniquely identifies the IPspace.


|===

//end collapsible .Definitions block
====

:leveloffset: -1


:leveloffset: -1


= Object-store


:leveloffset: +1


= Manage S3 buckets


:leveloffset: +1


[[IDa78f6179ab595477f1c1d2998ddf0681]]
= Protocols S3 buckets endpoint overview


== Overview

```

An S3 bucket is a container of objects. Each bucket defines an object namespace. S3 server requests specify objects using a bucket-name and object-name pair. An object consists of data, along with optional metadata and access controls, that is accessible using a name. An object resides within a bucket. There can be more than one bucket in an S3 server. Buckets that are created for the server are associated with an S3 user that is created on the S3 server.

== Examples

=== Retrieving all fields for all S3 buckets of a cluster

The API:

/api/protocols/s3/buckets

The call:

```
curl -X GET "https://<mgmt-  
ip>/api/protocols/s3/buckets?fields=*&return_records=true" -H "accept:  
application/json"
```

The response:

```
{  
  "records": [  
    {  
      "svm": {  
        "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",  
        "name": "vs1"  
      },  
      "uuid": "527812ab-7c6d-11e9-97e8-0050568ea123",  
      "name": "bucket-2",  
      "volume": {  
        "name": "fg_oss_1558514455",  
        "uuid": "51276f5f-7c6d-11e9-97e8-0050568ea123"  
      },  
      "size": 209715200,  
      "logical_used_size": 157286400,  
      "encryption": {  
        "enabled": false  
      },  
      "comment": "S3 bucket.",  
      "qos_policy": {  
        "min_throughput_iops": 0,  
        "max_throughput_iops": 1000,  
        "max_throughput_mbps": 0,  
      }  
    }  
  ]  
}
```



```

    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
    "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
  }
},
{
  "svm": {
    "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
    "name": "vs1"
  },
  "uuid": "a8234aec-7e06-11e9-97e8-0050568ea123",
  "name": "bucket-1",
  "volume": {
    "name": "fg_oss_1558690256",
    "uuid": "a36a1ea7-7e06-11e9-97e8-0050568ea123"
  },
  "size": 1677721600,
  "logical_used_size": 0,
  "encryption": {
    "enabled": false
  },
  "comment": "bucket2",
  "qos_policy": {
    "min_throughput_iops": 0,
    "max_throughput_iops": 1000,
    "max_throughput_mbps": 0,
    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
    "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
  }
},
{
  "svm": {
    "uuid": "ee30eb2d-7ae1-11e9-8abe-0050568ea123",
    "name": "vs2"
  },
  "uuid": "19283b75-7ae2-11e9-8abe-0050568ea123",
  "name": "bucket-3",
  "volume": {
    "name": "fg_oss_1558690257",
    "uuid": "a46a1ea7-7e06-11e9-97e8-0050568ea123"
  },
  "size": 1677721600,
  "logical_used_size": 1075838976,
  "encryption": {
    "enabled": false
  },
  "comment": "bucket3",

```

```

"qos_policy": {
  "min_throughput_iops": 0,
  "max_throughput_iops": 1000,
  "max_throughput_mbps": 0,
  "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
  "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
},
"policy": {
  "statements": [
    {
      "effect": "allow",
      "actions": [
        "*"
      ],
      "principals": [
        "Alice"
      ],
      "resources": [
        "bucket-3",
        "bucket-3/*"
      ],
      "sid": "fullAccessForAliceToBucket",
      "conditions": [
        {
          "operator": "ip_address",
          "source_ips": [
            "1.1.1.1/10"
          ]
        }
      ]
    }
  ]
}
}
],
"num_records": 3
}
-----

```

=== Retrieving all S3 buckets of a cluster ordered by size

```

# The API:
/api/protocols/s3/buckets

```

```

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/s3/buckets?return_records=true&order_by=size" -H
"accept: application/json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
        "name": "vs1"
      },
      "uuid": "754389d0-7e13-11e9-bfdc-0050568ea123",
      "name": "bb1",
      "size": 83886080
    },
    {
      "svm": {
        "uuid": "ee30eb2d-7ae1-11e9-8abe-0050568ea123",
        "name": "vs2"
      },
      "uuid": "19283b75-7ae2-11e9-8abe-0050568ea123",
      "name": "bb2",
      "size": 838860800
    },
    {
      "svm": {
        "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
        "name": "vs1"
      },
      "uuid": "a8234aec-7e06-11e9-97e8-0050568ea123",
      "name": "bucket-1",
      "size": 1677721600
    }
  ],
  "num_records": 3
}
----

=== Retrieving all S3 buckets of a cluster with name "bb2"

----

# The API:
/api/protocols/s3/buckets

```

```
# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/s3/buckets?name=bb2&return_records=true" -H "accept:
application/json"
```

```
# The response:
```

```
{
"records": [
  {
    "svm": {
      "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
      "name": "vs1"
    },
    "uuid": "087d940e-7e15-11e9-bfdc-0050568ea123",
    "name": "bb2"
  },
  {
    "svm": {
      "uuid": "ee30eb2d-7ae1-11e9-8abe-0050568ea123",
      "name": "vs2"
    },
    "uuid": "19283b75-7ae2-11e9-8abe-0050568ea123",
    "name": "bb2"
  }
],
"num_records": 2
}
----
```

```
=== Retrieving the specified bucket associated with an SVM
```

```
----
```

```
# The API:
```

```
/api/protocols/s3/buckets/{svm.uuid}/{uuid}
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/s3/buckets/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/527812ab-7c6d-11e9-97e8-0050568ea123" -H "accept:
application/json"
```

```
# The response:
```

```
{
"svm": {
  "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
```

```

    "name": "vs1"
  },
  "uuid": "527812ab-7c6d-11e9-97e8-0050568ea123",
  "name": "bucket-2",
  "volume": {
    "name": "fg_oss_1558514455",
    "uuid": "51276f5f-7c6d-11e9-97e8-0050568ea123"
  },
  "size": 209715200,
  "logical_used_size": 157286400,
  "encryption": {
    "enabled": false
  },
  "comment": "S3 bucket.",
  "qos_policy": {
    "min_throughput_iops": 0,
    "max_throughput_iops": 1000,
    "max_throughput_mbps": 0,
    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
    "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
  }
}

```

=== Creating an S3 bucket for an SVM

The API:

/api/protocols/s3/buckets

The call:

```

curl -iku admin:netapp1! -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
  \"aggregates\": [ { \"name\": \"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-
8c06-0050568ea123\" } ], \"comment\": \"S3 bucket.\",
  \"constituents_per_aggregate\": 4, \"name\": \"bucket-3\", \"svm\": {
    \"name\": \"vs1\" } }"

```

The response:

HTTP/1.1 202 Accepted

Date: Fri, 24 May 2019 11:22:14 GMT

Server: libzapid-httpd

X-Content-Type-Options: nosniff

Cache-Control: no-cache,no-store,must-revalidate

```

Location: /api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "name": "bucket-3",
      "comment": "S3 bucket."
    }
  ],
  "job": {
    "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
      }
    }
  }
}
-----

```

=== Creating an S3 bucket along with QoS policy for an SVM

```

-----

# The API:
/api/protocols/s3/buckets

# The call:
curl -iku admin:netapp1! -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
  \"comment\": \"S3 bucket.\", \"name\": \"bucket-3\", \"svm\": { \"name\":
  \"vs1\" }, \"qos_policy\": { \"min_throughput_iops\": 0,
  \"max_throughput_iops\": 1000000, \"max_throughput_mbps\": 900000,
  \"uuid\": \"02d07a93-6177-11ea-b241-000c293feac8\", \"name\":
  \"vs0_auto_gen_policy_02cfa02a_6177_11ea_b241_000c293feac8\" } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate

```

```

Location: /api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "name": "bucket-3",
      "comment": "S3 bucket."
    }
  ],
  "job": {
    "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
      }
    }
  }
}
-----

```

=== Creating an S3 bucket along with policies and conditions for an SVM

The API:

/api/protocols/s3/buckets

The call:

```

curl -iku admin:netappl! -X POST "https://<mgmt-
ip>/api/protocols/s3/buckets?return_timeout=0&return_records=true" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
  \"aggregates\": [ { \"name\": \"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-
8c06-0050568ea123\" } ], \"comment\": \"S3 bucket.\",
  \"constituents_per_aggregate\": 4, \"name\": \"bucket-3\", \"policy\": {
  \"statements\": [ { \"actions\": [ \"GetObject\" ], \"conditions\": [ {
  \"operator\": \"ip_address\", \"source_ips\": [ \"1.1.1.1/23\",
  \"1.2.2.2/20\" ] } ], \"effect\": \"allow\", \"resources\": [ \"bucket-
3/policies/examples/*\" ], \"sid\": \"AccessToGetObjectForAllUsersofSVM\"
  }, { \"actions\": [ \"*Object\" ], \"effect\": \"deny\", \"principals\": [
  \"mike\" ], \"resources\": [ \"bucket-3/policy-docs/*\", \"bucket-
3/confidential-*\" ], \"sid\": \"DenyAccessToObjectForMike\" }, {
  \"actions\": [ \"GetObject\" ], \"effect\": \"allow\", \"principals\": [
  \"*\" ], \"resources\": [ \"bucket-3/readme\" ], \"sid\":

```

```
\\"AnononymousAccessToGetObjectForUsers\\" } ] }, \\"svm\\": { \\"uuid\\":  
\\"259b4e46-2d33-11ea-9145-005056bbb1\\" } }"
```

The response:

HTTP/1.1 202 Accepted

Date: Fri, 24 May 2019 11:22:14 GMT

Server: libzapid-httpd

X-Content-Type-Options: nosniff

Cache-Control: no-cache,no-store,must-revalidate

Location: /api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbb1/?name=bucket-3

Content-Length: 353

Content-Type: application/json

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "name": "bucket-3",  
      "comment": "S3 bucket."  
    }  
  ],  
  "job": {  
    "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",  
    "_links": {  
      "self": {  
        "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"  
      }  
    }  
  }  
}
```

=== Updating an S3 bucket for an SVM

The API:

/api/protocols/s3/buckets/{svm.uuid}/{uuid}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/buckets/259b4e46-2d33-  
11ea-9145-005056bbb1/376a2efd-2d4d-11ea-9c30-  
005056bb883a?return_records=true" -H "accept:  
application/json?return_records=true" -H "Content-Type: application/json"  
-d "{ \\"comment\\": \\"Bucket modified.\\", \\"size\\": 11111111111,  
\\"qos_policy\\": { \\"min_throughput_iops\\": 0, \\"max_throughput_iops\\":
```



```
1000000, \"max_throughput_mbps\": 900000, \"uuid\": \"02d07a93-6177-11ea-b241-000c293feac8\", \"name\":  
\"vs0_auto_gen_policy_02cfa02a_6177_11ea_b241_000c293feac8\" } }"
```

The response:

```
HTTP/1.1 202 Accepted  
Date: Fri, 24 May 2019 11:32:27 GMT  
Server: libzapid-httpd  
X-Content-Type-Options: nosniff  
Cache-Control: no-cache,no-store,must-revalidate  
Content-Length: 189  
Content-Type: application/json
```

```
{  
  "job": {  
    "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",  
    "_links": {  
      "self": {  
        "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"  
      }  
    }  
  }  
}
```

=== Updating an S3 bucket policy for an SVM

The API:

```
/api/protocols/s3/buckets/{svm.uuid}/{uuid}
```

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-005056bbb883a?return_records=true" -H "accept: application/json" -H  
"Content-Type: application/json" -d "{ \"policy\": { \"statements\": [ {  
  \"actions\": [ \"*\"] , \"conditions\": [ { \"operator\": \"ip_address\",  
  \"source_ips\": [ \"1.1.1.5/23\" ] } ] , \"effect\": \"allow\",  
  \"resources\": [ \"*\"] , \"sid\": \"fullAccessForAllPrincipalsToBucket\" }  
] } }"
```

The response:

```
HTTP/1.1 202 Accepted  
Date: Fri, 24 May 2019 11:32:27 GMT  
Server: libzapid-httpd  
X-Content-Type-Options: nosniff
```

```
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
  "job": {
    "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
      }
    }
  }
}
```

=== Deleting an S3 bucket for a specified SVM

```
----

# The API:
/api/protocols/s3/buckets/{svm.uuid}/{uuid}

# The call:
curl -iku admin:netapp1! -X DELETE "https://<mgmt-
ip>/api/protocols/s3/buckets/259b4e46-2d33-11ea-9145-
005056bbbec1/98528221-2d52-11ea-892e-005056bbbec1?return_records=true" -H
"accept: application/json"
```

```
# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:40:17 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
  "job": {
    "uuid": "b3af4a54-7e18-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/b3af4a54-7e18-11e9-bfdc-0050568ea123"
      }
    }
  }
}
```

[[IDf8aca58c4d5675675ea81cb79c3c014a]]
= Retrieve all S3 buckets for all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/s3/buckets`#

Introduced In: 9.7

Retrieves all S3 buckets for all SVMs. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '**'.

== Related ONTAP commands

- * `vserver object-store-server bucket show`
- * `vserver object-store-server bucket policy statement show`
- * `vserver object-store-server bucket policy-statement-condition show`

== Learn more

* xref:{relative_path}protocols_s3_buckets_endpoint_overview.html[DOC /protocols/s3/buckets]

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

size
integer
query
False
a Filter by size

```
|volume.uuid
|string
|query
|False
a|Filter by volume.uuid
```

```
|volume.name
|string
|query
|False
a|Filter by volume.name
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|policy.statements.conditions.source_ips
|string
|query
|False
a|Filter by policy.statements.conditions.source_ips
```

* Introduced in: 9.8

```
|policy.statements.conditions.prefixes
|string
|query
```

```
|False
a|Filter by policy.statements.conditions.prefixes

* Introduced in: 9.8
```

```
|policy.statements.conditions.operator
|string
|query
|False
a|Filter by policy.statements.conditions.operator

* Introduced in: 9.8
```

```
|policy.statements.conditions.max_keys
|integer
|query
|False
a|Filter by policy.statements.conditions.max_keys

* Introduced in: 9.8
```

```
|policy.statements.conditions.usernames
|string
|query
|False
a|Filter by policy.statements.conditions.usernames

* Introduced in: 9.8
```

```
|policy.statements.conditions.delimiters
|string
|query
|False
a|Filter by policy.statements.conditions.delimiters

* Introduced in: 9.8
```

```
|policy.statements.principals
|string
|query
|False
a|Filter by policy.statements.principals
```

* Introduced in: 9.8

```
|policy.statements.sid  
|string  
|query  
|False  
a|Filter by policy.statements.sid
```

* Introduced in: 9.8

```
|policy.statements.effect  
|string  
|query  
|False  
a|Filter by policy.statements.effect
```

* Introduced in: 9.8

```
|policy.statements.actions  
|string  
|query  
|False  
a|Filter by policy.statements.actions
```

* Introduced in: 9.8

```
|policy.statements.resources  
|string  
|query  
|False  
a|Filter by policy.statements.resources
```

* Introduced in: 9.8

```
|logical_used_size  
|integer  
|query  
|False  
a|Filter by logical_used_size
```

```

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid


|svm.name
|string
|query
|False
a|Filter by svm.name


|qos_policy.max_throughput_mbps
|integer
|query
|False
a|Filter by qos_policy.max_throughput_mbps


* Introduced in: 9.8


|qos_policy.max_throughput_iops
|integer
|query
|False
a|Filter by qos_policy.max_throughput_iops


* Introduced in: 9.8


|qos_policy.uuid
|string
|query
|False
a|Filter by qos_policy.uuid


* Introduced in: 9.8


|qos_policy.min_throughput_iops
|integer
|query
|False
a|Filter by qos_policy.min_throughput_iops

```

* Introduced in: 9.8

|qos_policy.min_throughput_mbps
|integer
|query
|False
a|Filter by qos_policy.min_throughput_mbps

* Introduced in: 9.8

|qos_policy.name
|string
|query
|False
a|Filter by qos_policy.name

* Introduced in: 9.8

|encryption.enabled
|boolean
|query
|False
a|Filter by encryption.enabled

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number

of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records

```
|records
|array[link:#s3_bucket[s3_bucket]]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "comment": "S3 bucket.",
      "logical_used_size": 0,
      "name": "bucket1",
      "policy": {
        "statements": [
          {
            "actions": [
              "GetObject",
              "PutObject",
              "DeleteObject",
              "ListBucket"
            ],
            "conditions": [
              {
                "delimiters": [
                  "/"
                ],
                "max_keys": [
                  "1000"
                ],
                "operator": "ip_address",
                "prefixes": [
                  "pref"
                ]
              }
            ]
          }
        ]
      }
    }
  ]
}
```

```

        ],
        "source_ips": [
            "1.1.1.1",
            "1.2.2.0/24"
        ],
        "usernames": [
            "user1"
        ]
    }
],
"effect": "allow",
"principals": [
    "user1",
    "group/grp1"
],
"resources": [
    "bucket1",
    "bucket1/*"
],
"sid": "FullAccessToUser1"
}
]
},
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},

```

```

    "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",

```

```

    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

```

```
[#encryption]
[.api-collapsible-fifth-title]
encryption
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|enabled
|boolean
a|Specifies whether encryption is enabled on the bucket. By default,
encryption is disabled on a bucket.
```

```
|===
```

```
[#s3_bucket_policy_condition]
[.api-collapsible-fifth-title]
s3_bucket_policy_condition
```

Information about policy conditions based on various condition operators and condition keys.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|delimiters
|array[string]
a|An array of delimiters that are compared with the delimiter value
specified at the time of execution of an S3-based command, using the
condition operator specified.
```

```
|max_keys
|array[integer]
a|An array of maximum keys that are allowed or denied to be retrieved
using an S3 list operation, based on the condition operator specified.
```

|operator
|string
a|Condition operator that is applied to the specified condition key.

|prefixes
|array[string]
a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips
|array[string]
a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames
|array[string]
a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]
[.api-collapsible-fifth-title]
s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|actions
|array[string]
a|

|conditions


```
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
```

```
a|Specifies bucket policy conditions.
```

```
|effect
```

```
|string
```

```
a|Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.
```

```
|principals
```

```
|array[string]
```

```
a|
```

```
|resources
```

```
|array[string]
```

```
a|
```

```
|sid
```

```
|string
```

```
a|Specifies the statement identifier used to differentiate between statements.
```

```
|===
```

```
[#policy]
```

```
[.api-collapsible-fifth-title]
```

```
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

```
|===
```

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|max_throughput_iops
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
```

```
|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
```

```
|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting
```

"min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_mbps

|integer

a|Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

|name

|string

a|The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.

|uuid

|string

a|The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the SVM.

|uuid

|string

```
a|The unique identifier of the SVM.
```

```
|===
```

```
[#volume]  
[.api-collapsible-fifth-title]  
volume
```

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|name  
|string  
a|The name of the volume.
```

```
|uuid  
|string  
a|Unique identifier for the volume. This corresponds to the instance-uuid  
that is exposed in the CLI and ONTAPI. It does not change due to a volume  
move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7  
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket]  
[.api-collapsible-fifth-title]  
s3_bucket
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name

pair. An object resides within a bucket.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|comment

|string

a|Can contain any additional information about the bucket being created or modified.

|encryption

|link:#encryption[encryption]

a|

|logical_used_size

|integer

a|Specifies the bucket logical used size up to this point.

|name

|string

a|Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-".

|policy

|link:#policy[policy]

a|A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

|qos_policy

|link:#qos_policy[qos_policy]

a|Specifes "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifes "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by

specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.
```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|Specifies the unique identifier of the bucket.
```

```
|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====

[[ID7964a4638142a9aea6a58ab037b26483]]
= Create the S3 bucket configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/s3/buckets`#

```

Introduced In: 9.8

Creates the S3 bucket configuration of an SVM.

== Important notes

- * Each SVM can have one or more bucket configurations.
- * Aggregate lists should be specified explicitly. If not specified, then the bucket is auto-provisioned as a FlexGroup volume.
- * Constituents per aggregate specifies the number of components (or FlexVol volumes) per aggregate. Is specified only when an aggregate list is explicitly defined.
- * An access policy can be created along with a bucket create. If creating an access policy fails, bucket configurations are saved and the access policy can be created using the PATCH endpoint.
- * "qos_policy" can be specified if a bucket needs to be attached to a QoS group policy during creation time.

== Required properties

- * `svm.uuid or svm.name` - Existing SVM in which to create the bucket configuration.
- * `name` - Bucket name that is to be created.

== Recommended optional properties

- * `aggregates` - List of aggregates for the FlexGroup volume on which the bucket is hosted on.
- * `constituents_per_aggregate` - Number of constituents per aggregate.
- * `size` - Specifying the bucket size is recommended.
- * `policy` - Specifying a policy enables users to perform operations on buckets; specifying the resource permissions is recommended.
- * `qos_policy` - A QoS policy for buckets.

== Default property values

- * `size` - 800MB
- * `comment` - ""
- * `aggregates` - No default value.
- * `constituents_per_aggregate` - _4_ , if an aggregates list is specified. Otherwise, no default value.
- * `policy.statements.actions` - GetObject, PutObject, DeleteObject, ListBucket.
- * `policy.statements.principals` - all S3 users and groups in the SVM.
- * `policy.statements.resources` - all objects in the bucket.
- * `policy.statements.conditions` - list of bucket policy conditions.

== Related ONTAP commands

- * ``vserver object-store-server bucket create``
- * ``vserver object-store-server bucket policy statement create``

== Learn more

- * `xref:{relative_path}protocols_s3_buckets_endpoint_overview.html` [DOC /protocols/s3/buckets]

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

return_timeout
integer
query
False

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

- * Default value: 1
- * Max value: 120
- * Min value: 0

return_records
boolean
query
False

a|The default is false. If set to true, the records are returned.

- * Default value:

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|aggregates
```

```
|array[link:#aggregates[aggregates]]
```

a|A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume.

```
|comment
```

```
|string
```

a|Can contain any additional information about the bucket being created or modified.

```
|constituents_per_aggregate
```

```
|integer
```

a|Specifies the number of constituents or FlexVol volumes per aggregate. A FlexGroup volume consisting of all such constituents across all specified aggregates is created. This option is used along with the aggregates option and cannot be used independently.

```
|encryption
```

```
|link:#encryption[encryption]
```

```
a|
```

```
|logical_used_size
```

```
|integer
```

a|Specifies the bucket logical used size up to this point.

```
|name
```

```
|string
```

a|Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-".

```
|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.
```

```
|qos_policy
|link:#qos_policy[qos_policy]
a|Specifies "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifies "min_throughput_iops" is only supported on volumes hosted on a
node that is flash optimized. A pre-created QoS policy can also be used by
specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or
assigning a QoS policy to a bucket is not supported if its containing
volume or SVM already has a QoS policy attached.
```

```
|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.
```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|Specifies the unique identifier of the bucket.
```

```
|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.
```

```
|===
```

```
.Example request
[%collapsible%closed]
```

```

=====
[source,json,subs=+macros]
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "comment": "S3 bucket.",
  "constituents_per_aggregate": "4",
  "logical_used_size": 0,
  "name": "bucket1",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              "1000"
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",

```

```

        "group/grp1"
    ],
    "resources": [
        "bucket1",
        "bucket1/*"
    ],
    "sid": "FullAccessToUser1"
}
]
},
"qos_policy": {
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {
    "name": "volumel",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Response

```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
=====

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error code | Message
//end header
//end row
//start row
|92405777 +
//end row
//start row
|"Failed to create bucket "{bucket name}" for SVM "{svm.name}". Reason:
{Reason of failure}. ";
//end row
//start row
|92405785 +

```

```

//end row
//start row
|"Bucket name "{bucket name}" contains invalid characters. Valid
characters for a bucket name are 0-9, a-z, ".", and "-". ";
//end row
//start row
|92405786 +
//end row
//start row
|"Bucket name "{bucket name}" is not valid. Bucket names must have between
3 and 63 characters. ";
//end row
//start row
|92405811 +
//end row
//start row
|"Failed to create bucket "{bucket name}" for SVM "{svm.name}". Wait a few
minutes and try the operation again.";
//end row
//start row
|92405812 +
//end row
//start row
|"Failed to create the object store volume. Reason: {Reason for
failure}.";
//end row
//start row
|92405819 +
//end row
//start row
|"Cannot provision an object store server volume for bucket "{bucket
name}" in SVM "{svm.name}" on the following aggregates because they are
SnapLock aggregates: {List of aggregates.name}.";
//end row
//start row
|92405820 +
//end row
//start row
|"Failed to check whether the aggregate "{aggregates.name}" is a
FabricPool. Reason: {Reason for failure}.";
//end row
//start row
|92405821 +
//end row
//start row
|"Cannot provision an object store server volume for bucket "{bucket

```

```

name}" in SVM "{svm.name}" on the following aggregates because they are
FabricPool: {List of aggregates.name}.";
//end row
//start row
|92405827 +
//end row
//start row
|"Internal Error. Unable to generate object store volume name.";
//end row
//start row
|92405857 +
//end row
//start row
|"One or more aggregates must be specified if "constituents_per_aggregate"
is specified.";
//end row
//start row
|92405858 +
//end row
//start row
|"Failed to "create" the "bucket" because the operation is only supported
on data SVMs.";
//end row
//start row
|92405859 +
//end row
//start row
|"The specified "aggregates.uuid" "{aggregates.uuid}" does not exist.";
//end row
//start row
|92405860 +
//end row
//start row
|"The specified "aggregates.name" "{aggregates.name}" and
"aggregates.uuid" "{aggregates.uuid}" refer to different aggregates.";
//end row
//start row
|92405861 +
//end row
//start row
|"The specified SVM UUID or bucket UUID does not exist.";
//end row
//start row
|92405863 +
//end row
//start row

```



```

|"An error occurs when creating an access policy. The reason for failure
is detailed in the error message.";
//end row
//start row
|92405891 +
//end row
//start row
|The resources specified in the access policy are not valid. Valid ways to
specify a resource are *, <bucket-name>, <bucket-name>/.../.... Valid
characters for a resource are 0-9, A-Z, a-z, _, +, comma, ;, :, =, ., &,
@,?, (, ), single quote, *, !, - and $.

//end row
//start row
|92405894 +
//end row
//start row
|"Statements, principals and resources list can have a maximum of 10
entries.";
//end row
//start row
|92405897 +
//end row
//start row
|The principals specified in the access policy are not in the correct
format. User name must be in between 1 and 64 characters. Valid characters
for a user name are 0-9, A-Z, a-z, _, +, =, comma, ., @, and - .

//end row
//start row
|92405898 +
//end row
//start row
|"The SID specified in the access policy is not valid. Valid characters
for a SID are 0-9, A-Z and a-z.";
//end row
|===
//end table

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#encryption]
[.api-collapsible-fifth-title]

```

encryption

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket.

|===

[#s3_bucket_policy_condition]

[.api-collapsible-fifth-title]

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|delimiters

|array[string]

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

|max_keys

|array[integer]

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

|operator

|string

a|Condition operator that is applied to the specified condition key.

|prefixes
|array[string]
a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips
|array[string]
a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames
|array[string]
a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]
[.api-collapsible-fifth-title]
s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|actions
|array[string]
a|

|conditions
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
a|Specifies bucket policy conditions.

```
|effect
|string
a|Specifies whether access is allowed or denied when a user requests the
specific action. If access (to allow) is not granted explicitly to a
resource, access is implicitly denied. Access can also be denied
explicitly to a resource, in order to make sure that a user cannot access
it, even if a different policy grants access.
```

```
|principals
|array[string]
a|
```

```
|resources
|array[string]
a|
```

```
|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.
```

```
|===
```

```
[#policy]
[.api-collapsible-fifth-title]
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

|===

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|max_throughput_iops
|integer
```

a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|max_throughput_mbps
|integer
```

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_iops
|integer
```

a|Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_mbps
|integer
```

a|Specifies the minimum throughput in Megabytes per sec, 0 means none.

This is mutually exclusive with name and UUID during POST and PATCH.

```
|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.
```

```
|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.
```

|===

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

|===

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket]
```

```
[.api-collapsible-fifth-title]
```

```
s3_bucket
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|aggregates
```

```
|array[link:#aggregates[aggregates]]
```

```
a|A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume.
```



```

|comment
|string
a|Can contain any additional information about the bucket being created or
modified.

|constituents_per_aggregate
|integer
a|Specifies the number of constituents or FlexVol volumes per aggregate. A
FlexGroup volume consisting of all such constituents across all specified
aggregates is created. This option is used along with the aggregates
option and cannot be used independently.

|encryption
|link:#encryption[encryption]
a|

|logical_used_size
|integer
a|Specifies the bucket logical used size up to this point.

|name
|string
a|Specifies the name of the bucket. Bucket name is a string that can only
contain the following combination of ASCII-range alphanumeric characters
0-9, a-z, ".", and "-".

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifes "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifes "min_throughput_iops" is only supported on volumes hosted on a
node that is flash optimized. A pre-created QoS policy can also be used by
specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or

```

assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

|size

|integer

a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|Specifies the unique identifier of the bucket.

|volume

|link:#volume[volume]

a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

|===

[#job_link]

[.api-collapsible-fifth-title]

job_link

[cols=3*,options=header]

|===

|Name

|Type

|Description

|uuid

|string

a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDd4b4b0b1fd0846eb2f623ed960f9596e]]
= Delete the S3 bucket configuration for an SVM

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/s3/buckets/{svm.uuid}/{uuid}`#

*Introduced In:* 9.8

Deletes the S3 bucket configuration of an SVM. An access policy is also
deleted on an S3 bucket "delete" command.

== Related ONTAP commands

* `vserver object-store-server bucket delete`
* `vserver object-store-server bucket policy statement delete`
* `vserver object-store-server bucket policy-statement-condition delete`

== Learn more

* xref:{relative_path}protocols_s3_buckets_endpoint_overview.html[DOC
/protocols/s3/buckets]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid

```

```

|string
|path
|True
a|The unique identifier of the bucket.

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds. This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job. If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.

* Default value: 1
* Max value: 120
* Min value: 0

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response

```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error code | Message
//end header
//end row
//start row
|92405811 +
//end row
//start row

```

```

|"Failed to delete bucket "{bucket name}" for SVM "{svm.name}". Wait a few
minutes and try the operation again.";
//end row
//start row
|92405858 +
//end row
//start row
|"Failed to "delete" the "bucket" because the operation is only supported
on data SVMs.";
//end row
//start row
|92405861 +
//end row
//start row
|"The specified SVM UUID or bucket UUID does not exist.";
//end row
//start row
|92405779 +
//end row
//start row
|"Failed to remove bucket "{bucket name}" for SVM "{svm.name}". Reason:
{Reason for failure}.";
//end row
//start row
|92405813 +
//end row
//start row
|"Failed to delete the object store volume. Reason: {Reason for
failure}.";
//end row
//start row
|92405864 +
//end row
//start row
|"An error occurred when deleting an access policy. The reason for failure
is detailed in the error message.";
//end row
|===
//end table

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href

```



```

|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]

```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID3f664a1981d454de282cab18ff613a59]]
= Retrieve the S3 bucket configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/s3/buckets/{svm.uuid}/{uuid}`#
```

Introduced In: 9.8

Retrieves the S3 bucket configuration of an SVM. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '***'.

== Related ONTAP commands

```
* `vserver object-store-server bucket show`
* `vserver object-store-server bucket policy statement show`
* `vserver object-store-server bucket policy-statement-condition show`
```

== Learn more

```
* xref:{relative_path}protocols_s3_buckets_endpoint_overview.html[DOC
/protocols/s3/buckets]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|uuid
|string
|path
|True
a|The unique identifier of the bucket.

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Can contain any additional information about the bucket being created or
modified.

|encryption
|link:#encryption[encryption]
a|

|logical_used_size
|integer
a|Specifies the bucket logical used size up to this point.
```

```

|name
|string
a|Specifies the name of the bucket. Bucket name is a string that can only
contain the following combination of ASCII-range alphanumeric characters
0-9, a-z, ".", and "-".

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifies "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifies "min_throughput_iops" is only supported on volumes hosted on a
node that is flash optimized. A pre-created QoS policy can also be used by
specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or
assigning a QoS policy to a bucket is not supported if its containing
volume or SVM already has a QoS policy attached.

|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|svm
|link:#svm[svm]
a|

|uuid
|string
a|Specifies the unique identifier of the bucket.

|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "comment": "S3 bucket.",
  "logical_used_size": 0,
  "name": "bucket1",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              "1000"
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",
```

```

        "group/grp1"
    ],
    "resources": [
        "bucket1",
        "bucket1/*"
    ],
    "sid": "FullAccessToUser1"
}
]
},
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "volumel",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
=====

```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```



```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#encryption]
[.api-collapsible-fifth-title]
encryption
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|enabled
|boolean
```

a|Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket.

```
|===
```

```
[#s3_bucket_policy_condition]
[.api-collapsible-fifth-title]
s3_bucket_policy_condition
```

Information about policy conditions based on various condition operators and condition keys.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|delimiters
|array[string]
```

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

|max_keys

|array[integer]

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

|operator

|string

a|Condition operator that is applied to the specified condition key.

|prefixes

|array[string]

a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips

|array[string]

a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames

|array[string]

a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]

[.api-collapsible-fifth-title]

s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

|actions
|array[string]
a|

|conditions
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
a|Specifies bucket policy conditions.

|effect
|string
a|Specifies whether access is allowed or denied when a user requests the
specific action. If access (to allow) is not granted explicitly to a
resource, access is implicitly denied. Access can also be denied
explicitly to a resource, in order to make sure that a user cannot access
it, even if a different policy grants access.

|principals
|array[string]
a|

|resources
|array[string]
a|

|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.

```

```

|===

```

```

[#policy]
[.api-collapsible-fifth-title]
policy

```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must

be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

```
|===
```

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|max_throughput_iops
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
```

```

|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.


|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH.


|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.


|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.


|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.


|===


[#svm]
[.api-collapsible-fifth-title]
svm


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links

```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the SVM.
```

```
|===
```

```
[#volume]
```

```
[.api-collapsible-fifth-title]
```

```
volume
```

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

|===

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```



```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDf96376e99e1ca31625e9d32da0905f90]]
= Update the S3 bucket configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/s3/buckets/{svm.uuid}/{uuid}`#

*Introduced In:* 9.8

Updates the S3 bucket configuration of an SVM.

== Important notes

* The following fields can be modified for a bucket:

*** `comment` - Any information related to the bucket.

*** `size` - Bucket size.

*** `policy` - An access policy for resources (buckets and objects) that
defines their permissions. New policies are created after existing
policies are deleted. To retain any of the existing policy statements, you
need to specify those statements again. Also, policy conditions can be
specified as part of a bucket policy.

*** `qos_policy` - A QoS policy for buckets.

== Related ONTAP commands

* `vserver object-store-server bucket modify`

```

```
* `vserver object-store-server bucket policy statement modify`  
* `vserver object-store-server bucket policy-statement-condition modify`
```

== Learn more

```
* xref:{relative_path}protocols_s3_buckets_endpoint_overview.html[DOC  
/protocols/s3/buckets]
```

== Parameters

```
[cols=5*,options=header]  
|==
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|uuid  
|string  
|path  
|True  
a|The unique identifier of the bucket.
```

```
|return_timeout  
|integer  
|query  
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

```
* Default value: 1  
* Max value: 120  
* Min value: 0
```

```
|svm.uuid  
|string  
|path
```

```

|True
a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Can contain any additional information about the bucket being created or
modified.

|encryption
|link:#encryption[encryption]
a|

|logical_used_size
|integer
a|Specifies the bucket logical used size up to this point.

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifes "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifes "min_throughput_iops" is only supported on volumes hosted on a
node that is flash optimized. A pre-created QoS policy can also be used by
specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or

```

assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

|size

|integer

a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|uuid

|string

a|Specifies the unique identifier of the bucket.

|volume

|link:#volume[volume]

a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "comment": "S3 bucket.",
  "logical_used_size": 0,
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              "1000"
            ],
            "operator": "ip_address",
```

```

        "prefixes": [
            "pref"
        ],
        "source_ips": [
            "1.1.1.1",
            "1.2.2.0/24"
        ],
        "usernames": [
            "user1"
        ]
    }
],
"effect": "allow",
"principals": [
    "user1",
    "group/grp1"
],
"resources": [
    "bucket1",
    "bucket1/*"
],
"sid": "FullAccessToUser1"
}
],
},
"qos_policy": {
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Response

```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
=====

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error code | Message
//end header
//end row
//start row
|92405778 +
//end row
//start row
|"Failed to modify bucket "{bucket name}" for SVM "{svm.name}". Reason:
{Reason for failure}."
//end row
//start row
|92405846 +

```

```

//end row
//start row
|"Failed to modify the object store volume. Reason: {Reason for
failure}.";
//end row
//start row
|92405811 +
//end row
//start row
|"Failed to modify bucket "{bucket name}" for SVM "{svm.name}". Wait a few
minutes and try the operation again.";
//end row
//start row
|92405858 +
//end row
//start row
|"Failed to "modify" the "bucket" because the operation is only supported
on data SVMs.";
//end row
//start row
|92405861 +
//end row
//start row
|"The specified SVM UUID or bucket UUID does not exist.";
//end row
//start row
|92405863 +
//end row
//start row
|"An error occurs when creating an access policy. The reason for failure
is detailed in the error message.";
//end row
//start row
|92405864 +
//end row
//start row
|"An error occurs when deleting an access policy. The reason for failure
is detailed in the error message.";
//end row
//start row
|92405891 +
//end row
//start row
|The resources specified in the access policy are not valid. Valid ways to
specify a resource are *, <bucket-name>, <bucket-name>/.../.... Valid
characters for a resource are 0-9, A-Z, a-z, _, +, comma, ;, :, =, ., &,

```

@,?, (,), single quote, *, !, - and \$.

```
//end row
//start row
|92405894 +
//end row
//start row
|"Statements, principals and resources list can have a maximum of 10
entries.";
//end row
//start row
|92405897 +
//end row
//start row
|The principals specified in the access policy are not in the correct
format. User name must be in between 1 and 64 characters. Valid characters
for a user name are 0-9, A-Z, a-z, _, +, =, comma, ., @, and - .

//end row
//start row
|92405898 +
//end row
//start row
|"The SID specified in the access policy is not valid. Valid characters
for a SID are 0-9, A-Z and a-z.";
//end row
|===
//end table
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```



```

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#encryption]
[.api-collapsible-fifth-title]
encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether encryption is enabled on the bucket. By default,

```

encryption is disabled on a bucket.

|===

```
[#s3_bucket_policy_condition]
[.api-collapsible-fifth-title]
s3_bucket_policy_condition
```

Information about policy conditions based on various condition operators and condition keys.

```
[cols=3*,options=header]
```

|===

|Name

|Type

|Description

|delimiters

|array[string]

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

|max_keys

|array[integer]

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

|operator

|string

a|Condition operator that is applied to the specified condition key.

|prefixes

|array[string]

a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips

|array[string]

a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames

|array[string]

a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]

[.api-collapsible-fifth-title]

s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|actions

|array[string]

a|

|conditions

|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]

a|Specifies bucket policy conditions.

|effect

|string

a|Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.

|principals

|array[string]

```

a|

|resources
|array[string]
a|

|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.

```

```

|===

```

```

[#policy]
[.api-collapsible-fifth-title]
policy

```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```

[cols=3*,options=header]

```

```

|===

```

```

|Name
|Type
|Description

```

```

|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.

```

```

|===

```

```

[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy

```

Specifes "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".

Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|max_throughput_iops
```

```
|integer
```

a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|max_throughput_mbps
```

```
|integer
```

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_iops
```

```
|integer
```

a|Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_mbps
```

```
|integer
```

a|Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|name
```

```
|string
```

a|The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.

```
|uuid
```

```
|string
```

a|The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#volume]

[.api-collapsible-fifth-title]

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the volume.

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket]
[.api-collapsible-fifth-title]
s3_bucket
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|comment
|string
a|Can contain any additional information about the bucket being created or
modified.
```

```
|encryption
|link:#encryption[encryption]
a|
```

```
|logical_used_size
|integer
a|Specifies the bucket logical used size up to this point.
```

```
|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
```

(bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|uuid
|string
a|Specifies the unique identifier of the bucket.

|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid


```

|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code

```

```

|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

```

|===

```

```

//end collapsible .Definitions block
====

```

```

:leveloffset: -1

```

```

= Manage protocols S3 services

```

```

:leveloffset: +1

```

```

[[IDac87942fd2d6fda3297f08d928194473]]
= Protocols S3 services endpoint overview

```

```

== Overview

```

An S3 server is an object store server that is compatible with the Amazon S3 protocol. In the initial version, only a subset of the protocol features necessary to support Fabric Pool capacity tier usecases are implemented. S3 server allows you to store objects in ONTAP using Amazon S3 protocol. This feature can be used as a target object store server for ONTAP FabricPools.

```

== Performance monitoring

```

Performance of the SVM can be monitored by the `metric.+++` and

`statistics.++` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The `metric.++` properties denote an average whereas `statistics.++` properties denote a real-time monotonically increasing value aggregated across all nodes.

== Examples

=== Retrieving all of the S3 configurations

The API:

/api/protocols/s3/services

The call:

```
curl -X GET "https://<mgmt-  
ip>/api/protocols/s3/services?fields=*&return_records=true&return_timeout=  
15" -H "accept: application/json"
```

The response:

```
{  
  "records": [  
    {  
      "svm": {  
        "uuid": "cf90b8f2-8071-11e9-8190-0050568eae21",  
        "name": "vs2"  
      },  
      "name": "s1",  
      "comment": "S3 server",  
      "enabled": false,  
    },  
    {  
      "svm": {  
        "uuid": "d7f1219c-7f8e-11e9-9124-0050568eae21",  
        "name": "vs1"  
      },  
      "name": "Server-1",  
      "comment": "S3 server",  
      "enabled": true,  
      "buckets": [  
        {  
          "uuid": "e08665af-8114-11e9-8190-0050568eae21",  
          "name": "bucket-1",  
          "volume": {  
            "name": "fg_oss_1559026220",  
            "uuid": "de146bfff-8114-11e9-8190-0050568eae21"  
          }  
        }  
      ]  
    }  
  ]  
}
```

```

    },
    "size": 209715200,
    "logical_used_size": 157286400,
    "encryption": {
        "enabled": false
    },
    "comment": "s3 bucket"
},
{
    "uuid": "fb1912ef-8114-11e9-8190-0050568eae21",
    "name": "bucket-2",
    "volume": {
        "name": "fg_oss_1559026269",
        "uuid": "f9b1cdd0-8114-11e9-8190-0050568eae21"
    },
    "size": 104857600,
    "logical_used_size": 78643200,
    "encryption": {
        "enabled": false
    },
    "comment": "s3 bucket"
}
],
"users": [
    {
        "name": "user-1",
        "comment": "S3 user",
        "access_key":
"3333_wl62ypaTi7_aAQuJo76Z16zc9Gz_W3IN83bDQWkcCN3jYU_z_xn20XATMKKa90509KCH
__r4lh1lIPU58vf1QlWAJt8k2F1BPjPtM6CsDRX_dOP_QZkF5N9fBuz3"
    },
    {
        "name": "user-2",
        "comment": "",
        "access_key":
"g6T24qhH92dOA6gc1WTcDO_2oNZhQ6Dr12zu5_s5Id_QK1wLgghgxSD2xPlxqG7oX1T_9AI0D
39q65CY3FAg0CbAtVU_903bSnCnht3xqjbrF5_3Cs9RnY8nE_az1Ltc"
    }
]
}
],
"num_records": 2
}
----

```

=== Retrieving all S3 configurations for a particular SVM

The API:

/api/protocols/s3/services/{svm.uuid}

The call:

```
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/24c2567a-f269-11e8-8852-0050568e5298?fields=*" -H "accept: application/json"
```

The response:

```
{
  "svm": {
    "uuid": "d7f1219c-7f8e-11e9-9124-0050568eae21",
    "name": "vs1"
  },
  "name": "Server-1",
  "comment": "S3 server",
  "enabled": true,
  "buckets": [
    {
      "uuid": "e08665af-8114-11e9-8190-0050568eae21",
      "name": "bucket-1",
      "volume": {
        "name": "fg_oss_1559026220",
        "uuid": "de146bff-8114-11e9-8190-0050568eae21"
      },
      "size": 209715200,
      "logical_used_size": 157286400,
      "encryption": {
        "enabled": false
      },
      "comment": "s3 bucket",
      "policy": {
        "statements": [
          {
            "effect": "deny",
            "actions": [
              "*Object"
            ],
            "principals": [
              "mike"
            ],
            "resources": [
              "bucket-1/policy-docs/*",
              "bucket-1/confidential-*"
            ]
          }
        ]
      }
    }
  ]
}
```

```

    ],
    "sid": "DenyAccessToGetPutDeleteObjectForMike"
  },
  {
    "effect": "allow",
    "actions": [
      "GetObject"
    ],
    "principals": [
      "*"
    ],
    "resources": [
      "bucket-1/readme"
    ],
    "sid": "AccessToGetObjectForAnonymousUser"
  }
]
}
},
{
  "uuid": "fb1912ef-8114-11e9-8190-0050568eae21",
  "name": "bucket-2",
  "volume": {
    "name": "fg_oss_1559026269",
    "uuid": "f9b1cdd0-8114-11e9-8190-0050568eae21"
  },
  "size": 1677721600,
  "logical_used_size": 1075838976,
  "encryption": {
    "enabled": false
  },
  "comment": "s3 bucket"
}
],
"users": [
  {
    "name": "user-1",
    "comment": "s3 user",
    "access_key":
"3333_wl62ypaTi7_aAQuJo76Zl6zc9Gz_W3IN83bDQWkcCN3jYU_z_xn20XATMKKa90509KCH
__r4lh1IPU58vf1QlWAJt8k2F1BPjPtM6CsDRX_dOP_QZkF5N9fBuz3"
  },
  {
    "name": "user-2",
    "comment": "",
    "access_key":

```

```

"6T24qhH92dOA6gc1WTcDO_2oNZhQ6Drl2zu5_s5Id_QK1wLgghgxsD2xPlxqG7oX1T_9AI0D
39q65CY3FAg0CbAtVU_903bSnCnht3xqjbrF5_3Cs9RnY8nE_az1Ltc"
    }
]
}
----

=== Creating an S3 server, users, and buckets configurations with required
fields specified

----

# The API:
/api/protocols/s3/services

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"buckets\": [ { \"name\": \"bucket-1\" }, { \"name\": \"bucket-2\" } ],
\"enabled\": true, \"name\": \"Server-1\", \"svm\": { \"uuid\":
\"d49ef663-7f8e-11e9-9b2c-0050568e4594\" }, \"users\": [ { \"name\":
\"user-1\" }, { \"name\": \"user-2\" } ]}"

# The response:
HTTP/1.1 201 Created
Date: Fri, 31 May 2019 08:44:16 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/
Content-Length: 623
Content-Type: application/hal+json
{
  "num_records": 1,
  "records": [
    {
      "users": [
        {
          "name": "user-1",
          "access_key":
"x129aL0q9bu3J_4_2S0OcU34AA5DJXXB_j9R34_60tqiqAS5_c8PAgN6Lg1zkv_76P4IxNWIr
9st9uhhgldb31u364Cczq_c39C1fUP7HDheUmYY6u4xt61_N7Sw6c33",
          "secret_key":
"gh0pYc__43Csnx_Ks4_C0tb_5AfT4HZTfQl8xN8Dl5TjqB90oNt5ZaPO6Hs4h6Q4Fq4B4uq5C
qht82X6vcE32c3uLZB8pXAAx819LWPgpOSwD5xga2RE3czrlqhCd9V6"
        },

```

```

    {
      "name": "user-2",
      "access_key":
"nntYZrNN65mKn57yS04o1sDp_D0AY58jdwCW573_5x2OPW09AbyF186DB7r30N2373_bA12n0
8aovQp8ySItRss9AjsYoSj7TsIiHOW_Y21DaqYP15I2a849b11y8X4c",
      "secret_key":
"bjtsPXV2D8BM6pZnQ9pzmKoXU3qIv2yQ3957dhjK4X7M2dB6Rjtrq1As_8cS_4bSP0jt_P31R
5eLdZ_zcBO9Z_ZRMldTclBw_5c7LugBnzG2D3xxB91jqLaP2xnKn_Zg"
    }
  ],
  "job": {
    "uuid": "f51675dd-820a-11e9-a762-0050568e4594",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f51675dd-820a-11e9-a762-0050568e4594"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/protocols/s3/services/"
    }
  }
}
]
}

```

=== Creating an S3 server, users, and buckets configurations

The API:

/api/protocols/s3/services

The call:

```

curl -X POST "https://<mgmt-ip>/api/protocols/s3/services" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
  \"buckets\": [ { \"aggregates\": [ { \"name\": \"aggr1\", \"uuid\":
  \"1cd8a442-86d1-11e0-ae1c-123478563412\" } ],
  \"constituents_per_aggregate\": 4, \"name\": \"bucket-1\", \"size\":
  \"209715200\", \"policy\": { \"statements\": [ { \"actions\": [ \"*\"] },
  \"conditions\": [ { \"operator\": \"ip_address\", \"source_ips\": [
  \"1.1.1.1/23\", \"1.2.2.2/20\" ] } ], \"effect\": \"allow\",
  \"resources\": [ \"bucket-1\", \"bucket-1*\" ], \"sid\":
  \"fullAccessForAllPrincipalsToBucket\" } ] } }, { \"aggregates\": [ {

```



```
\ "name\": \"aggr1\", \"uuid\": \"1cd8a442-86d1-11e0-ae1c-123478563412\" },
{ \"name\": \"aggr2\", \"uuid\": \"982fc4d0-d1a2-4da4-9c47-5b433f24757d\"}
], \"constituents_per_aggregate\": 4, \"name\": \"bucket-2\" } ],
\"enabled\": true, \"name\": \"Server-1\", \"svm\": { \"name\": \"vs1\",
\"uuid\": \"d49ef663-7f8e-11e9-9b2c-0050568e4594\" }, \"users\": [ {
\"name\": \"user-1\" }, { \"name\": \"user-2\" } ]}"
```

The response:

HTTP/1.1 201 Created

Date: Fri, 31 May 2019 08:44:16 GMT

Server: libzapid-httpd

X-Content-Type-Options: nosniff

Cache-Control: no-cache,no-store,must-revalidate

Location: /api/protocols/s3/services/

Content-Length: 623

Content-Type: application/hal+json

```
{
  "num_records": 1,
  "records": [
    {
      "users": [
        {
          "name": "user-1",
          "access_key":
"x129aL0q9bu3J_4_2S0OcU34AA5DJXXB_j9R34_60tqiqAS5_c8PAgN6Lglzkv_76P4IxnWir
9st9uhhgldb31u364Cczq_c39C1fUP7HDheUmYY6u4xt61_N7Sw6c33",
          "secret_key":
"gh0pYc__43Csnx_Ks4_C0tb_5AfT4HZTfQl8xN8Dl5TjqB90oNt5ZaPO6Hs4h6Q4Fq4B4uq5C
qht82X6vcE32c3uLZB8pXAAx819LWPgpOSwD5xga2RE3czrlqhCd9V6"
        },
        {
          "name": "user-2",
          "access_key":
"nntYZrNN65mKn57yS04o1sDp_D0AY58jdwCW573_5x2OPW09AbyF186DB7r30N2373_bA12n0
8aovQp8ySItRss9AjsYoSj7TsIiHOW_Y21DaqYP15I2a849b11y8X4c",
          "secret_key":
"bjtsPXV2D8BM6pZnQ9pzmKoXU3qIv2yQ3957dhjK4X7M2dB6Rjtrq1As_8cS_4bSP0jt_P31R
5eLdZ_zcBO9Z_ZRMldTc1Bw_5c7LugBnzG2D3xXB91jqLaP2xnKn_Zg"
        }
      ],
      "job": {
        "uuid": "f51675dd-820a-11e9-a762-0050568e4594",
        "_links": {
          "self": {
            "href": "/api/cluster/jobs/f51675dd-820a-11e9-a762-0050568e4594"
          }
        }
      }
    }
  ]
}
```

```

    }
  },
  "_links": {
    "self": {
      "href": "/api/protocols/s3/services/"
    }
  }
}
]
}
----

```

=== Creating an S3 server configuration

The API:

/api/protocols/s3/services

The call:

```

curl -X POST "https://<mgmt-ip>/api/protocols/s3/services" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"comment\":
\"S3 server\", \"enabled\": true, \"name\": \"Server-1\", \"svm\": {
\"name\": \"vs1\", \"uuid\": \"db2ec036-8375-11e9-99e1-0050568e3ed9\" } }"
----

```

=== Disable s3 server for the specified SVM

The API:

/api/protocols/s3/services/{svm.uuid}

The call:

```

curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/03ce5c36-f269-
11e8-8852-0050568e5298" -H "accept: application/json" -H "Content-Type:
application/json" -d "{ \"enabled\": false }"
----

```

=== Deleting the S3 server for a specified SVM

The API:

/api/protocols/s3/services/{svm.uuid}

The call:

```

curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/a425f10b-ad3b-
11e9-b559-0050568e8222?delete_all=false" -H "accept: application/json"
HTTP/1.1 200 OK
Date: Wed, 14 Aug 2019 07:04:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 132
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "job": {
        "uuid": "bf74ba50-be61-11e9-bea8-0050568e8222"
      }
    }
  ]
}
-----

=== Deleting all of the S3 server configuration for a specified SVM

-----

# The API:
/api/protocols/s3/services/{svm.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/03ce5c36-f269-
11e8-8852-0050568e5298?delete_all=true" -H "accept: application/json"

# The response:
HTTP/1.1 200 OK
Date: Sat, 01 Jun 2019 15:46:39 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 132
Content-Type: application/hal+json
{
  "num_records": 1,
  "records": [
    {
      "job": {
        "uuid": "71eaaf02-8484-11e9-91f7-0050568ebc5f"
      }
    }
  ]
}

```

```

    }
  }
]
}
----

```

[[IDa09e6a824cf3b72648279b5b8ff6ed9d]]
 = Retrieve the S3 server configuration for all SVMs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/s3/services`#

Introduced In: 9.7

Retrieves the S3 server configuration for all SVMs. Note that in order to retrieve S3 bucket policy conditions, 'fields' option should be set to '**'.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `statistics.+++`
 * `metric.+++`

== Related ONTAP commands

* `vserver object-store-server show`

== Learn more

* [xref:{relative_path}protocols_s3_services_endpoint_overview.html\[DOC /protocols/s3/services\]](#)

== Parameters

[cols=5*,options=header]
 |===

```
|Name
|Type
|In
|Required
|Description

|secure_port
|integer
|query
|False
a|Filter by secure_port
```

* Introduced in: 9.8

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|buckets.size
|integer
|query
|False
a|Filter by buckets.size
```

```
|buckets.volume.uuid
|string
|query
|False
a|Filter by buckets.volume.uuid
```

```
|buckets.volume.name
|string
|query
```

```

|False
a|Filter by buckets.volume.name

|buckets.uuid
|string
|query
|False
a|Filter by buckets.uuid

|buckets.comment
|string
|query
|False
a|Filter by buckets.comment

|buckets.name
|string
|query
|False
a|Filter by buckets.name

|buckets.policy.statements.conditions.source_ips
|string
|query
|False
a|Filter by buckets.policy.statements.conditions.source_ips

* Introduced in: 9.8

|buckets.policy.statements.conditions.prefixes
|string
|query
|False
a|Filter by buckets.policy.statements.conditions.prefixes

* Introduced in: 9.8

|buckets.policy.statements.conditions.operator
|string
|query
|False

```

a|Filter by buckets.policy.statements.conditions.operator

* Introduced in: 9.8

|buckets.policy.statements.conditions.max_keys

|integer

|query

|False

a|Filter by buckets.policy.statements.conditions.max_keys

* Introduced in: 9.8

|buckets.policy.statements.conditions.usernames

|string

|query

|False

a|Filter by buckets.policy.statements.conditions.usernames

* Introduced in: 9.8

|buckets.policy.statements.conditions.delimiters

|string

|query

|False

a|Filter by buckets.policy.statements.conditions.delimiters

* Introduced in: 9.8

|buckets.policy.statements.principals

|string

|query

|False

a|Filter by buckets.policy.statements.principals

* Introduced in: 9.8

|buckets.policy.statements.sid

|string

|query

|False

a|Filter by buckets.policy.statements.sid

* Introduced in: 9.8

```
|buckets.policy.statements.effect
|string
|query
|False
a|Filter by buckets.policy.statements.effect
```

* Introduced in: 9.8

```
|buckets.policy.statements.actions
|string
|query
|False
a|Filter by buckets.policy.statements.actions
```

* Introduced in: 9.8

```
|buckets.policy.statements.resources
|string
|query
|False
a|Filter by buckets.policy.statements.resources
```

* Introduced in: 9.8

```
|buckets.logical_used_size
|integer
|query
|False
a|Filter by buckets.logical_used_size
```

```
|buckets.svm.uuid
|string
|query
|False
a|Filter by buckets.svm.uuid
```

```
|buckets.svm.name
|string
|query
```



```

|False
a|Filter by buckets.svm.name

|buckets.qos_policy.max_throughput_mbps
|integer
|query
|False
a|Filter by buckets.qos_policy.max_throughput_mbps

* Introduced in: 9.8

|buckets.qos_policy.max_throughput_iops
|integer
|query
|False
a|Filter by buckets.qos_policy.max_throughput_iops

* Introduced in: 9.8

|buckets.qos_policy.uuid
|string
|query
|False
a|Filter by buckets.qos_policy.uuid

* Introduced in: 9.8

|buckets.qos_policy.min_throughput_iops
|integer
|query
|False
a|Filter by buckets.qos_policy.min_throughput_iops

* Introduced in: 9.8

|buckets.qos_policy.min_throughput_mbps
|integer
|query
|False
a|Filter by buckets.qos_policy.min_throughput_mbps

* Introduced in: 9.8

```

```
|buckets.qos_policy.name
|string
|query
|False
a|Filter by buckets.qos_policy.name
```

* Introduced in: 9.8

```
|buckets.encryption.enabled
|boolean
|query
|False
a|Filter by buckets.encryption.enabled
```

```
|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total
```

* Introduced in: 9.8

```
|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read
```

* Introduced in: 9.8

```
|statistics.latency_raw.other
|integer
|query
|False
a|Filter by statistics.latency_raw.other
```

* Introduced in: 9.8

```
|statistics.latency_raw.write
|integer
```

```
|query
|False
a|Filter by statistics.latency_raw.write
```

* Introduced in: 9.8

```
|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp
```

* Introduced in: 9.8

```
|statistics.iops_raw.total
|integer
|query
|False
a|Filter by statistics.iops_raw.total
```

* Introduced in: 9.8

```
|statistics.iops_raw.read
|integer
|query
|False
a|Filter by statistics.iops_raw.read
```

* Introduced in: 9.8

```
|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other
```

* Introduced in: 9.8

```
|statistics.iops_raw.write
|integer
|query
|False
```

```

a|Filter by statistics.iops_raw.write

* Introduced in: 9.8

|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write

* Introduced in: 9.8

|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read

* Introduced in: 9.8

|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total

* Introduced in: 9.8

|statistics.status
|string
|query
|False
a|Filter by statistics.status

* Introduced in: 9.8

|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write

```

* Introduced in: 9.8

```
|metric.throughput.read  
|integer  
|query  
|False  
a|Filter by metric.throughput.read
```

* Introduced in: 9.8

```
|metric.throughput.total  
|integer  
|query  
|False  
a|Filter by metric.throughput.total
```

* Introduced in: 9.8

```
|metric.duration  
|string  
|query  
|False  
a|Filter by metric.duration
```

* Introduced in: 9.8

```
|metric.timestamp  
|string  
|query  
|False  
a|Filter by metric.timestamp
```

* Introduced in: 9.8

```
|metric.status  
|string  
|query  
|False  
a|Filter by metric.status
```

* Introduced in: 9.8

```
|metric.iops.total
|integer
|query
|False
a|Filter by metric.iops.total
```

* Introduced in: 9.8

```
|metric.iops.read
|integer
|query
|False
a|Filter by metric.iops.read
```

* Introduced in: 9.8

```
|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other
```

* Introduced in: 9.8

```
|metric.iops.write
|integer
|query
|False
a|Filter by metric.iops.write
```

* Introduced in: 9.8

```
|metric.latency.total
|integer
|query
|False
a|Filter by metric.latency.total
```

* Introduced in: 9.8

```
|metric.latency.read
```

```
|integer
|query
|False
a|Filter by metric.latency.read
```

* Introduced in: 9.8

```
|metric.latency.other
|integer
|query
|False
a|Filter by metric.latency.other
```

* Introduced in: 9.8

```
|metric.latency.write
|integer
|query
|False
a|Filter by metric.latency.write
```

* Introduced in: 9.8

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|certificate.name
|string
|query
|False
a|Filter by certificate.name
```

* Introduced in: 9.8

```
|certificate.uuid  
|string  
|query  
|False  
a|Filter by certificate.uuid
```

* Introduced in: 9.8

```
|enabled  
|boolean  
|query  
|False  
a|Filter by enabled
```

```
|port  
|integer  
|query  
|False  
a|Filter by port
```

* Introduced in: 9.8

```
|is_https_enabled  
|boolean  
|query  
|False  
a|Filter by is_https_enabled
```

* Introduced in: 9.8

```
|is_http_enabled  
|boolean  
|query  
|False  
a|Filter by is_http_enabled
```

* Introduced in: 9.8

```
|users.access_key  
|string
```



```
|query
|False
a|Filter by users.access_key
```

```
|users.svm.uuid
|string
|query
|False
a|Filter by users.svm.uuid
```

```
|users.svm.name
|string
|query
|False
a|Filter by users.svm.name
```

```
|users.comment
|string
|query
|False
a|Filter by users.comment
```

```
|users.name
|string
|query
|False
a|Filter by users.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```

|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

```

```
|num_records
|integer
a|Number of records
```

```
|records
|array[link:#s3_service[s3_service]]
a|

|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "buckets": [
        {
          "comment": "S3 bucket.",
          "logical_used_size": 0,
          "name": "bucket1",
          "policy": {
            "statements": [
              {
                "actions": [
                  "GetObject",
                  "PutObject",
                  "DeleteObject",
                  "ListBucket"
                ]
              }
            ]
          }
        }
      ]
    }
  ]
}
```

```

    ],
    "conditions": [
      {
        "delimiters": [
          "/"
        ],
        "max_keys": [
          "1000"
        ],
        "operator": "ip_address",
        "prefixes": [
          "pref"
        ],
        "source_ips": [
          "1.1.1.1",
          "1.2.2.0/24"
        ],
        "usernames": [
          "user1"
        ]
      }
    ],
    "effect": "allow",
    "principals": [
      "user1",
      "group/grp1"
    ],
    "resources": [
      "bucket1",
      "bucket1/*"
    ],
    "sid": "FullAccessToUser1"
  }
],
},
"qos_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"max_throughput_iops": "10000",
"max_throughput_mbps": "500",
"min_throughput_iops": "2000",
"min_throughput_mbps": "500",
"name": "performance",

```

```

        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "size": "1677721600",
    "svm": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
    "volume": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
}
],
"certificate": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"comment": "S3 server",
"metric": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "duration": "PT15S",
    "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
    }
},

```

```

    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "name": "Server-1",
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "access_key":

```

```

"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41

```

```

U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
  "comment": "S3 user",
  "name": "user-1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
]
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",

```

```

        "message": "string"
    }
],
"code": "4",
"message": "entry doesn't exist",
"target": "uuid"
}
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]

```



```

a|

|self
|link:#href[href]
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

```

Aggregate

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#encryption]
[.api-collapsible-fifth-title]
encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether encryption is enabled on the bucket. By default,
encryption is disabled on a bucket.

|===

[#s3_bucket_policy_condition]
[.api-collapsible-fifth-title]
s3_bucket_policy_condition
```

Information about policy conditions based on various condition operators and condition keys.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|delimiters

|array[string]

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

|max_keys

|array[integer]

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

|operator

|string

a|Condition operator that is applied to the specified condition key.

|prefixes

|array[string]

a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips

|array[string]

a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames

|array[string]

a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

```
[#s3_bucket_policy_statement]
[.api-collapsible-fifth-title]
s3_bucket_policy_statement
```

Specifies information about a single access permission.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|actions
|array[string]
a|
```

```
|conditions
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
a|Specifies bucket policy conditions.
```

```
|effect
|string
a|Specifies whether access is allowed or denied when a user requests the
specific action. If access (to allow) is not granted explicitly to a
resource, access is implicitly denied. Access can also be denied
explicitly to a resource, in order to make sure that a user cannot access
it, even if a different policy grants access.
```

```
|principals
|array[string]
a|
```

```
|resources
|array[string]
a|
```

```
|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.
```

|===

```
[#policy]
[.api-collapsible-fifth-title]
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

|===

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
|===
|Name
```

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|max_throughput_iops
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually
exclusive with name and UUID during POST and PATCH.

|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH.

|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.

|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.

|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.

|===

```

```
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
```

a|The name of the volume.

|uuid

|string

a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.

* example: 028baa66-41bd-11e9-81d5-00a0986138f7

* Introduced in: 9.6

|===

[#s3_bucket]

[.api-collapsible-fifth-title]

s3_bucket

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|comment

|string

a|Can contain any additional information about the bucket being created or modified.

|encryption

|link:#encryption[encryption]

a|

|logical_used_size

|integer

a|Specifies the bucket logical used size up to this point.

|name


```

|string
a|Specifies the name of the bucket. Bucket name is a string that can only
contain the following combination of ASCII-range alphanumeric characters
0-9, a-z, ".", and "-".

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifies "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifies "min_throughput_iops" is only supported on volumes hosted on a
node that is flash optimized. A pre-created QoS policy can also be used by
specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or
assigning a QoS policy to a bucket is not supported if its containing
volume or SVM already has a QoS policy attached.

|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|svm
|link:#svm[svm]
a|

|uuid
|string
a|Specifies the unique identifier of the bucket.

|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

|===

```

```
[#certificate]
[.api-collapsible-fifth-title]
certificate
```

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
```

```
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

```

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]  
[.api-collapsible-fifth-title]  
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]  
[.api-collapsible-fifth-title]  
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```


a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw

|link:#latency_raw[latency_raw]

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#s3_user]

[.api-collapsible-fifth-title]

s3_user

This is a container of S3 users.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|access_key
```

```
|string
```

```
a|Specifies the access key for the user.
```

```
|comment
```

```
|string
```

```
a|Can contain any additional information about the user being created or modified.
```

```
|name
```

```
|string
```

```
a|Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

```
[#s3_service]
```

```
[.api-collapsible-fifth-title]
```

```
s3_service
```

```
Specifies the S3 server configuration.
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```

|link:#self_link[self_link]
a|

|buckets
|array[link:#s3_bucket[s3_bucket]]
a|

|certificate
|link:#certificate[certificate]
a|Specifies the certificate that will be used for creating HTTPS
connections to the S3 server.

|comment
|string
a|Can contain any additional information about the server being created or
modified.

|enabled
|boolean
a|Specifies whether the S3 server being created or modified should be up
or down.

|is_http_enabled
|boolean
a|Specifies whether HTTP is enabled on the S3 server being created or
modified. By default, HTTP is disabled on the S3 server.

|is_https_enabled
|boolean
a|Specifies whether HTTPS is enabled on the S3 server being created or
modified. By default, HTTPS is enabled on the S3 server.

|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput, for SVM
protocols.

|name
|string
a|Specifies the name of the S3 server. A server name can contain 0 to 15
characters using only the following combination of characters':' 0-9, A-Z,

```

a-z, ".", and "-".

|port

|integer

a|Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80.

|secure_port

|integer

a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443.

|statistics

|link:#statistics[statistics]

a|These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

|svm

|link:#svm[svm]

a|

|users

|array[link:#s3_user[s3_user]]

a|

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

[[ID361e612584d86b02325830de44b1d20f]]

= Create S3 server, users, and buckets configurations

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/protocols/s3/services`#

Introduced In: 9.7

Creates an S3 server, users, and buckets configurations.

== Important notes

- * Each SVM can have one S3 server configuration.
- * One or more buckets and users can also be created using this end-point.
- * If creating a user configuration fails, buckets are not created either and already created users are not saved.
- * If creating a bucket configuration fails, all buckets already created are saved with no new buckets created.

== Required properties

- * `svm.uuid` - Existing SVM in which to create an S3 server configuration.

== Recommended optional properties

- * `enabled` - Specifies the state of the server created.
- * `comment` - Any information related to the server created.

== Default property values

- * `comment` - ""
- * `enabled` - `_true_`

== Related ONTAP commands

- * `vserver object-store-server create`
- * `vserver object-store-server bucket create`
- * `vserver object-store-server bucket policy statement create`
- * `vserver object-store-server bucket policy-statement-condition create`
- * `vserver object-store-server user create`

== Learn more

- * `xref:{relative_path}protocols_s3_services_endpoint_overview.html` [DOC /protocols/s3/services]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===

|Name
|Type
|Description

|buckets
|array[link:#s3_bucket[s3_bucket]]

a|

|certificate
|link:#certificate[certificate]

a|Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

|comment
|string

a|Can contain any additional information about the server being created or modified.

```

|enabled
|boolean
a|Specifies whether the S3 server being created or modified should be up
or down.

|is_http_enabled
|boolean
a|Specifies whether HTTP is enabled on the S3 server being created or
modified. By default, HTTP is disabled on the S3 server.

|is_https_enabled
|boolean
a|Specifies whether HTTPS is enabled on the S3 server being created or
modified. By default, HTTPS is enabled on the S3 server.

|name
|string
a|Specifies the name of the S3 server. A server name can contain 0 to 15
characters using only the following combination of characters': ' 0-9, A-Z,
a-z, ".", and "-".

|port
|integer
a|Specifies the HTTP listener port for the S3 server. By default, HTTP is
enabled on port 80.

|secure_port
|integer
a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS
is enabled on port 443.

|svm
|link:#svm[svm]
a|

|users
|array[link:#s3_user[s3_user]]
a|

|===

```


.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "buckets": [
    {
      "aggregates": [
        {
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "comment": "S3 bucket.",
      "constituents_per_aggregate": "4",
      "logical_used_size": 0,
      "name": "bucket1",
      "policy": {
        "statements": [
          {
            "actions": [
              "GetObject",
              "PutObject",
              "DeleteObject",
              "ListBucket"
            ],
            "conditions": [
              {
                "delimiters": [
                  "/"
                ],
                "max_keys": [
                  "1000"
                ],
                "operator": "ip_address",
                "prefixes": [
                  "pref"
                ],
                "source_ips": [
                  "1.1.1.1",
                  "1.2.2.0/24"
                ],
                "usernames": [
                  "user1"
                ]
              }
            ]
          }
        ]
      }
    }
  ]
}
```

```

        }
    ],
    "effect": "allow",
    "principals": [
        "user1",
        "group/grp1"
    ],
    "resources": [
        "bucket1",
        "bucket1/*"
    ],
    "sid": "FullAccessToUser1"
}
]
},
"qos_policy": {
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
],
"certificate": {
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"comment": "S3 server",
"name": "Server-1",
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"users": [

```

```
{
  "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
  "comment": "S3 user",
  "name": "user-1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
]
}
=====
```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of Records


|records
|array[link:#records[records]]
a|

|===


.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "job": {
        "uuid": "string"
      },
      "users": [
        {
          "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
          "name": "user-1",
          "secret_key":
"A20_tDhC_cux2C2BmtL45bXB_a_Q65c_96FsAcOdo14Az8V31jBKDTc0uCL62Bh559gPB8s9r
rn0868QrF38_1dsV2u1_9H2tSf3qQ5xp9NT259C6z_GiZQ883Qn63X1"
        }
      ]
    }
  ]
}
=====

== Error

```

ONTAP Error Response Codes

|===

| Error Code | Description

| 2621706

| The specified SVM UUID is incorrect for the specified SVM name.

| 92405789

| The specified object server name contains invalid characters or not a fully qualified domain name. Valid characters for an object store server name are 0-9, A-Z, a-z, ".", and "-".

| 92405790

| Object store server names must have between 1 and 15 characters.

| 92405839

| Creating an object store server requires an effective cluster version of data ONTAP 9.7.0 or later. Upgrade all the nodes to 9.7.0 or later and try the operation again.

| 92405853

| Failed to create the object store server because Cloud Volumes ONTAP does not support object store servers.

| 92405863

| An error occurs when creating an S3 user or bucket. The reason for failure is detailed in the error message. Follow the error codes specified for the user or bucket endpoints to see details for the failure.

| 92405884

| An object store server can only be created on a data SVM. An object store server can also be created on a system SVM on a mixed platform cluster.

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

```

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link
[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

```

```

[#encryption]
[.api-collapsible-fifth-title]
encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether encryption is enabled on the bucket. By default,
encryption is disabled on a bucket.

|===

[#s3_bucket_policy_condition]
[.api-collapsible-fifth-title]
s3_bucket_policy_condition

Information about policy conditions based on various condition operators
and condition keys.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|delimiters
|array[string]
a|An array of delimiters that are compared with the delimiter value
specified at the time of execution of an S3-based command, using the
condition operator specified.

|max_keys
|array[integer]
a|An array of maximum keys that are allowed or denied to be retrieved
using an S3 list operation, based on the condition operator specified.

|operator

```

|string
a|Condition operator that is applied to the specified condition key.

|prefixes
|array[string]
a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips
|array[string]
a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames
|array[string]
a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]
[.api-collapsible-fifth-title]
s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|actions
|array[string]
a|

|conditions
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
a|Specifies bucket policy conditions.


```
|effect
|string
a|Specifies whether access is allowed or denied when a user requests the
specific action. If access (to allow) is not granted explicitly to a
resource, access is implicitly denied. Access can also be denied
explicitly to a resource, in order to make sure that a user cannot access
it, even if a different policy grants access.
```

```
|principals
|array[string]
a|
```

```
|resources
|array[string]
a|
```

```
|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.
```

```
|===
```

```
[#policy]
[.api-collapsible-fifth-title]
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|statements
```

```
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
```

```
a|Specifies bucket access policy statement.
```

```
|===
```

```
[#qos_policy]
```

```
[.api-collapsible-fifth-title]
```

```
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or

"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".

Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|max_throughput_iops
```

```
|integer
```

a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|max_throughput_mbps
```

```
|integer
```

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_iops
```

```
|integer
```

a|Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_mbps
```

```

|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.

|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.

|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#volume]
[.api-collapsible-fifth-title]
volume

```

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket]
```

```
[.api-collapsible-fifth-title]
```

```
s3_bucket
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|aggregates
```

```
|array[link:#aggregates[aggregates]]
```

```
a|A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume.
```

```

|comment
|string
a|Can contain any additional information about the bucket being created or
modified.

|constituents_per_aggregate
|integer
a|Specifies the number of constituents or FlexVol volumes per aggregate. A
FlexGroup volume consisting of all such constituents across all specified
aggregates is created. This option is used along with the aggregates
option and cannot be used independently.

|encryption
|link:#encryption[encryption]
a|

|logical_used_size
|integer
a|Specifies the bucket logical used size up to this point.

|name
|string
a|Specifies the name of the bucket. Bucket name is a string that can only
contain the following combination of ASCII-range alphanumeric characters
0-9, a-z, ".", and "-".

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifes "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifes "min_throughput_iops" is only supported on volumes hosted on a

```

node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.
```

```
|svm
|link:#svm[svm]
a|

|uuid
|string
a|Specifies the unique identifier of the bucket.
```

```
|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.
```

```
|===
```

```
[#certificate]
[.api-collapsible-fifth-title]
certificate
```

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
```

latency

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput]
```

```
[.api-collapsible-fifth-title]
```

throughput

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|duration
```

```
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
```

```
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
```

```
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#iops_raw]

[.api-collapsible-fifth-title]

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

```
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#s3_user]
[.api-collapsible-fifth-title]
s3_user

This is a container of S3 users.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|access_key
|string
a|Specifies the access key for the user.

|comment
|string
a|Can contain any additional information about the user being created or modified.

|name
|string
a|Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|svm
|link:#svm[svm]
a|

|===

```
[#s3_service]
[.api-collapsible-fifth-title]
s3_service
```

Specifies the S3 server configuration.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|buckets
|array[link:#s3_bucket[s3_bucket]]
a|
```

```
|certificate
|link:#certificate[certificate]
a|Specifies the certificate that will be used for creating HTTPS
connections to the S3 server.
```

```
|comment
|string
a|Can contain any additional information about the server being created or
modified.
```

```
|enabled
|boolean
a|Specifies whether the S3 server being created or modified should be up
or down.
```

```
|is_http_enabled
|boolean
a|Specifies whether HTTP is enabled on the S3 server being created or
modified. By default, HTTP is disabled on the S3 server.
```

```
|is_https_enabled
|boolean
```

a|Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server.

|name

|string

a|Specifies the name of the S3 server. A server name can contain 0 to 15 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-".

|port

|integer

a|Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80.

|secure_port

|integer

a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443.

|svm

|link:#svm[svm]

a|

|users

|array[link:#s3_user[s3_user]]

a|

|===

[#collection_links]

[.api-collapsible-fifth-title]

collection_links

[#job_link]

[.api-collapsible-fifth-title]

job_link

[cols=3*,options=header]

|===

|Name

|Type

|Description


```
|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.
```

```
|===
```

```
[#s3_service_user_post_response]
[.api-collapsible-fifth-title]
s3_service_user_post_response
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|access_key
```

```
|string
```

```
a|Specifies the access key for the user.
```

```
|name
```

```
|string
```

```
a|The name of the user.
```

```
|secret_key
```

```
|string
```

```
a|Specifies the secret key for the user.
```

```
|===
```

```
[#records]
[.api-collapsible-fifth-title]
records
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```

|job
|link:#job_link[job_link]
a|

|users
|array[link:#s3_service_user_post_response[s3_service_user_post_response]]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[IDdde434c99ea3bd11c2934905e726d2a9]]
= Delete the S3 server configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}`#
```

Introduced In: 9.7

Deletes the S3 server configuration of an SVM. If the 'delete_all' parameter is set to false, only the S3 server is deleted. Otherwise S3 users and buckets present on the SVM are also deleted. Note that only empty buckets can be deleted. This endpoint returns the S3 server delete job-uuid in response. To monitor the job status follow
/api/cluster/jobs/+++<job-uuid>+++.
+++</job-uuid>+++

== Related ONTAP commands

* `vserver object-store-server delete`

== Learn more

* xref:{relative_path}protocols_s3_services_endpoint_overview.html[DOC

```
/protocols/s3/services]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|delete_all
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|Delete S3 server and associated users and empty buckets.
```

```
* Default value: 1
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the SVM to which this object belongs.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of Records


|records
|array[link:#records[records]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "job": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "string"
      }
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 92405864

| An error occurs when deleting an S3 user or bucket. The reason for failure is detailed in the error message. Follow the error codes specified for the user or bucket endpoints to see details for the failure.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

====

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#job_link]
```

```
[.api-collapsible-fifth-title]
```

```
job_link
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```

|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#records]
[.api-collapsible-fifth-title]
records

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```



```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDe04c75228f2dc218715b12be237c1c07]]
= Retrieve the S3 server configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
```

```
block]#`/protocols/s3/services/{svm.uuid}`#
```

Introduced In: 9.7

Retrieves the S3 Server configuration of an SVM. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '*'.

== Related ONTAP commands

* `vserver object-store-server show`

== Learn more

* xref:{relative_path}protocols_s3_services_endpoint_overview.html[DOC /protocols/s3/services]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|==

== Response

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#self_link[self_link]
a|

|buckets
|array[link:#s3_bucket[s3_bucket]]
a|

|certificate
|link:#certificate[certificate]
a|Specifies the certificate that will be used for creating HTTPS
connections to the S3 server.

|comment
|string
a|Can contain any additional information about the server being created or
modified.

|enabled
|boolean
a|Specifies whether the S3 server being created or modified should be up
or down.

|is_http_enabled
|boolean
a|Specifies whether HTTP is enabled on the S3 server being created or
modified. By default, HTTP is disabled on the S3 server.

|is_https_enabled
|boolean
a|Specifies whether HTTPS is enabled on the S3 server being created or
modified. By default, HTTPS is enabled on the S3 server.
```

```

|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput, for SVM
protocols.

|name
|string
a|Specifies the name of the S3 server. A server name can contain 0 to 15
characters using only the following combination of characters': ' 0-9, A-Z,
a-z, ".", and "-".

|port
|integer
a|Specifies the HTTP listener port for the S3 server. By default, HTTP is
enabled on port 80.

|secure_port
|integer
a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS
is enabled on port 443.

|statistics
|link:#statistics[statistics]
a|These are raw performance numbers, such as IOPS latency and throughput
for SVM protocols. These numbers are aggregated across all nodes in the
cluster and increase with the uptime of the cluster.

|svm
|link:#svm[svm]
a|

|users
|array[link:#s3_user[s3_user]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]

```

```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "buckets": [
    {
      "comment": "S3 bucket.",
      "logical_used_size": 0,
      "name": "bucket1",
      "policy": {
        "statements": [
          {
            "actions": [
              "GetObject",
              "PutObject",
              "DeleteObject",
              "ListBucket"
            ],
            "conditions": [
              {
                "delimiters": [
                  "/"
                ],
                "max_keys": [
                  "1000"
                ],
                "operator": "ip_address",
                "prefixes": [
                  "pref"
                ],
                "source_ips": [
                  "1.1.1.1",
                  "1.2.2.0/24"
                ],
                "usernames": [
                  "user1"
                ]
              }
            ],
            "effect": "allow",
            "principals": [
              "user1",
              "group/grp1"
            ],

```

```

        "resources": [
            "bucket1",
            "bucket1/*"
        ],
        "sid": "FullAccessToUser1"
    }
]
},
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
},
"certificate": {
    "_links": {
        "self": {

```

```

        "href": "/api/resourcelink"
    },
    },
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"comment": "S3 server",
"metric": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    },
    "duration": "PT15S",
    "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    },
    "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"name": "Server-1",
"statistics": {
    "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    },
    "status": "ok",
    "throughput_raw": {

```

```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"users": [
    {
        "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
        "comment": "S3 user",
        "name": "user-1",
        "svm": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
    }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```
|error
|link:#error[error]
a|
```

```
|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
```

```

a|

|===

[#self_link]
[.api-collapsible-fifth-title]
self_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|name
|string
a|

```

```

|uuid
|string
a|

```

```

|===

```

```

[#encryption]
[.api-collapsible-fifth-title]
encryption

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|enabled
|boolean
a|Specifies whether encryption is enabled on the bucket. By default,
encryption is disabled on a bucket.

```

```

|===

```

```

[#s3_bucket_policy_condition]
[.api-collapsible-fifth-title]
s3_bucket_policy_condition

```

Information about policy conditions based on various condition operators and condition keys.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|delimiters
```

```
|array[string]
```

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

```
|max_keys
```

```
|array[integer]
```

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

```
|operator
```

```
|string
```

a|Condition operator that is applied to the specified condition key.

```
|prefixes
```

```
|array[string]
```

a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

```
|source_ips
```

```
|array[string]
```

a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

```
|usernames
```

```
|array[string]
```

a|An array of usernames that a current user in the context is evaluated against using the condition operators.

```
|===
```

```
[#s3_bucket_policy_statement]
[.api-collapsible-fifth-title]
s3_bucket_policy_statement
```

Specifies information about a single access permission.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|actions
```

```
|array[string]
```

```
a|
```

```
|conditions
```

```
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
```

```
a|Specifies bucket policy conditions.
```

```
|effect
```

```
|string
```

```
a|Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.
```

```
|principals
```

```
|array[string]
```

```
a|
```

```
|resources
```

```
|array[string]
```

```
a|
```

```
|sid
```

```
|string
```

```
a|Specifies the statement identifier used to differentiate between statements.
```

```
|===
```

```
[#policy]
[.api-collapsible-fifth-title]
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

```
|===
```

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
```

|link:#_links[_links]

a|

|max_throughput_iops

|integer

a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

|max_throughput_mbps

|integer

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_iops

|integer

a|Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_mbps

|integer

a|Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

|name

|string

a|The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.

|uuid

|string

a|The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume.
```

```
|uuid
```



```
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket]
[.api-collapsible-fifth-title]
s3_bucket
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|comment
```

```
|string
```

```
a|Can contain any additional information about the bucket being created or
modified.
```

```
|encryption
```

```
|link:#encryption[encryption]
```

```
a|
```

```
|logical_used_size
```

```
|integer
```

```
a|Specifies the bucket logical used size up to this point.
```

```
|name
```

```
|string
```

```
a|Specifies the name of the bucket. Bucket name is a string that can only
contain the following combination of ASCII-range alphanumeric characters
0-9, a-z, ".", and "-".
```

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|svm
|link:#svm[svm]
a|

|uuid
|string
a|Specifies the unique identifier of the bucket.

|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

|===

[#certificate]
[.api-collapsible-fifth-title]

certificate

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|Certificate name
```

```
|uuid
```

```
|string
```

```
a|Certificate UUID
```

```
|===
```

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS latency and throughput, for SVM protocols.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#statistics]
[.api-collapsible-fifth-title]
statistics

```

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

```

```
|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#s3_user]
[.api-collapsible-fifth-title]
s3_user
```

This is a container of S3 users.

```
[cols=3*,options=header]
|===
```

```

|Name
|Type
|Description

|access_key
|string
a|Specifies the access key for the user.

|comment
|string
a|Can contain any additional information about the user being created or
modified.

|name
|string
a|Specifies the name of the user. A user name length can range from 1 to
64 characters and can only contain the following combination of characters
0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID6fcdacbfdf0a19c656ed082e0855bad1]]
= Update the S3 server configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
```

```
block]#`/protocols/s3/services/{svm.uuid}`#
```

Introduced In: 9.7

Updates the S3 Server configuration of an SVM.

== Related ONTAP commands

* `vserver object-store-server modify`

== Learn more

* xref:{relative_path}protocols_s3_services_endpoint_overview.html[DOC
/protocols/s3/services]

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|svm.uuid  
|string  
|path  
|True  
a|UUID of the SVM to which this object belongs.
```

```
|===
```

== Request Body

```
[cols=3*,options=header]  
|===
```

```
|Name  
|Type  
|Description
```

```
|certificate  
|link:#certificate[certificate]
```

a|Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

|comment

|string

a|Can contain any additional information about the server being created or modified.

|enabled

|boolean

a|Specifies whether the S3 server being created or modified should be up or down.

|is_http_enabled

|boolean

a|Specifies whether HTTP is enabled on the S3 server being created or modified. By default, HTTP is disabled on the S3 server.

|is_https_enabled

|boolean

a|Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server.

|name

|string

a|Specifies the name of the S3 server. A server name can contain 0 to 15 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-".

|port

|integer

a|Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80.

|secure_port

|integer

a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443.

```
|===
```

```
.Example request
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "certificate": {
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "comment": "S3 server",
  "name": "Server-1"
}
=====
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 92405789
```

```
| The specified object server name contains invalid characters. Valid
characters for an object store server name are 0-9, A-Z, a-z, ".", and "-".
```

```
| 92405790
```

```
| Object store server names must have between 1 and 15 characters.
```

```
|===
```

```
== Definitions
```

```
[.api-def-first-level]
```

```
.See Definitions
```

```
[%collapsible%closed]
```

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#self_link]
```

```
[.api-collapsible-fifth-title]
```

```
self_link
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#aggregates]
```

```
[.api-collapsible-fifth-title]
```

```
aggregates
```

```
Aggregate
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```



```
[#encryption]
[.api-collapsible-fifth-title]
encryption
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

a|Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket.

```
|===
```

```
[#s3_bucket_policy_condition]
[.api-collapsible-fifth-title]
s3_bucket_policy_condition
```

Information about policy conditions based on various condition operators and condition keys.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|delimiters
```

```
|array[string]
```

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

```
|max_keys
```

```
|array[integer]
```

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

|operator
|string
a|Condition operator that is applied to the specified condition key.

|prefixes
|array[string]
a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips
|array[string]
a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames
|array[string]
a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]
[.api-collapsible-fifth-title]
s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|actions
|array[string]
a|

|conditions

```
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
```

```
a|Specifies bucket policy conditions.
```

```
|effect
```

```
|string
```

```
a|Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.
```

```
|principals
```

```
|array[string]
```

```
a|
```

```
|resources
```

```
|array[string]
```

```
a|
```

```
|sid
```

```
|string
```

```
a|Specifies the statement identifier used to differentiate between statements.
```

```
|===
```

```
[#policy]
```

```
[.api-collapsible-fifth-title]
```

```
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

```
|===
```

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|max_throughput_iops
```

```
|integer
```

a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|max_throughput_mbps
```

```
|integer
```

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_iops
```

```
|integer
```

a|Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.

|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.

|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#volume]
[.api-collapsible-fifth-title]

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket]
```

```
[.api-collapsible-fifth-title]
```

```
s3_bucket
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|comment
```

```
|string
```

```
a|Can contain any additional information about the bucket being created or
```

modified.

|encryption

|link:#encryption[encryption]

a|

|logical_used_size

|integer

a|Specifies the bucket logical used size up to this point.

|policy

|link:#policy[policy]

a|A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

|qos_policy

|link:#qos_policy[qos_policy]

a|Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

|size

|integer

a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|uuid

|string

a|Specifies the unique identifier of the bucket.

|volume

|link:#volume[volume]

a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
|===
```

```
[#certificate]  
[.api-collapsible-fifth-title]  
certificate
```

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|name  
|string  
a|Certificate name
```

```
|uuid  
|string  
a|Certificate UUID
```

```
|===
```

```
[#iops]  
[.api-collapsible-fifth-title]  
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|other  
|integer  
a|Performance metric for other I/O operations. Other I/O operations can be
```


metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#throughput]

[.api-collapsible-fifth-title]

throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#metric]

[.api-collapsible-fifth-title]

metric

Performance numbers, such as IOPS latency and throughput, for SVM

protocols.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|duration

|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops

|link:#iops[iops]

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

```
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
```

```
[.api-collapsible-fifth-title]
```

```
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

```
[.api-collapsible-fifth-title]
```

```
statistics
```

These are raw performance numbers, such as IOPS latency and throughput for SVM protocols. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

```
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
```

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
```

```
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp
```

```
|string
```

a|The timestamp of the performance data.

```
|===
```

```
[#s3_user]
```

```
[.api-collapsible-fifth-title]
```

```
s3_user
```

This is a container of S3 users.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

|Description

|access_key

|string

a|Specifies the access key for the user.

|comment

|string

a|Can contain any additional information about the user being created or modified.

|svm

|link:#svm[svm]

a|

|===

[#s3_service]

[.api-collapsible-fifth-title]

s3_service

Specifies the S3 server configuration.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|certificate

|link:#certificate[certificate]

a|Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

|comment

|string

a|Can contain any additional information about the server being created or modified.

|enabled

|boolean

a|Specifies whether the S3 server being created or modified should be up or down.

|is_http_enabled

|boolean

a|Specifies whether HTTP is enabled on the S3 server being created or modified. By default, HTTP is disabled on the S3 server.

|is_https_enabled

|boolean

a|Specifies whether HTTPS is enabled on the S3 server being created or modified. By default, HTTPS is enabled on the S3 server.

|name

|string

a|Specifies the name of the S3 server. A server name can contain 0 to 15 characters using only the following combination of characters':' 0-9, A-Z, a-z, ".", and "-".

|port

|integer

a|Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80.

|secure_port

|integer

a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage S3 service bucket configuration
```

```
:leveloffset: +1
```

```
[[ID9d3d4bf63676347ac4852fd8f64915a9]]
```

```
= Protocols S3 services svm.uuid buckets endpoint overview
```

```
== Overview
```

An S3 bucket is a container of objects. Each bucket defines an object namespace. S3 server requests specify objects using a bucket-name and object-name pair. An object consists of data, along with optional metadata and access controls, accessible via a name. An object resides within a bucket. There can be more than one bucket in an S3 server. Buckets which are created for the server are associated with an S3 user that is created on the S3 server.

An access policy is an object that when associated with a resource, defines their permissions. Buckets and objects are defined as resources. By default, only the "root" user can access these resources. Access policies are used to manage access to these resources by enabling ONTAP admin to provide "grants" to allow other users to perform operations on the buckets.

```
== Examples
```

```
=== Retrieving all fields for all S3 buckets of an SVM
```

```
----
```

```
# The API:
```

```
/api/protocols/s3/services/{svm.uuid}/buckets
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/buckets?fields=*&return_records=true" -H "accept: application/json"
```

```
# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
        "name": "vs1"
      },
      "uuid": "527812ab-7c6d-11e9-97e8-0050568ea123",
      "name": "bucket-2",
      "volume": {
        "name": "fg_oss_1558514455",
        "uuid": "51276f5f-7c6d-11e9-97e8-0050568ea123"
      },
      "size": 209715200,
      "logical_used_size": 157286400,
      "encryption": {
        "enabled": false
      },
      "comment": "S3 bucket.",
      "qos_policy": {
        "min_throughput_iops": 0,
        "max_throughput_iops": 1000,
        "max_throughput_mbps": 50,
        "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
        "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
      }
    },
    {
      "svm": {
        "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
        "name": "vs1"
      },
      "uuid": "a8234aec-7e06-11e9-97e8-0050568ea123",
      "name": "bucket-1",
      "volume": {
        "name": "fg_oss_1558690256",
        "uuid": "a36a1ea7-7e06-11e9-97e8-0050568ea123"
      },
      "size": 1677721600,
      "logical_used_size": 0,
      "encryption": {
        "enabled": false
      },
      "comment": "bucket1",

```

```

"qos_policy": {
  "min_throughput_iops": 0,
  "max_throughput_iops": 1000,
  "max_throughput_mbps": 50,
  "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
  "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
},
"policy": {
  "statements": [
    {
      "effect": "allow",
      "actions": [
        "*"
      ],
      "principals": [
        "Alice"
      ],
      "resources": [
        "*"
      ],
      "sid": "fullAccessForAliceToBucket"
    },
    {
      "effect": "allow",
      "actions": [
        "ListBucket",
        "GetObject"
      ],
      "principals": [
        "ann",
        "jack"
      ],
      "resources": [
        "bucket-1",
        "bucket-1/*"
      ],
      "sid": "AccessToListAndGetObjectForAnnAndJack",
      "conditions": [
        {
          "operator": "ip_address",
          "source_ips": [
            "1.1.1.1/10"
          ]
        },
        {
          "operator": "string_equals",

```

```

        "prefixes": [
            "pref1",
            "pref2"
        ],
        "usernames": [
            "user1",
            "user2"
        ],
        "delimiters": [
            "del1",
            "del2"
        ]
    },
    {
        "operator": "numeric_equals",
        "max_keys": [
            100
        ]
    }
]
},
{
    "effect": "deny",
    "actions": [
        "*Object"
    ],
    "principals": [
        "mike",
        "group/group1"
    ],
    "resources": [
        "bucket-1/policy-docs/*",
        "bucket-1/confidential-*"
    ],
    "sid": "DenyAccessToGetPutDeleteObjectForMike"
},
{
    "effect": "allow",
    "actions": [
        "GetObject"
    ],
    "principals": [
        "*"
    ],
    "resources": [
        "bucket-1/readme"
    ]
}

```

```

    ],
    "sid": "AccessToGetObjectForAnonymousUsers"
  },
  {
    "effect": "allow",
    "actions": [
      "GetObject"
    ],
    "principals": [
    ],
    "resources": [
      "bucket-1/policies/examples/*"
    ],
    "sid": "AccessToGetObjectForAllUsersOfSVM"
  }
]
}
}
],
"num_records": 2
}
----

```

=== Retrieving the specified bucket associated with an SVM

The API:

/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

The call:

```

curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/buckets/527812ab-7c6d-11e9-97e8-0050568ea123" -H
"accept: application/json"

```

The response:

```

{
  "svm": {
    "uuid": "12f3ba4c-7ae0-11e9-8c06-0050568ea123",
    "name": "vs1"
  },
  "uuid": "527812ab-7c6d-11e9-97e8-0050568ea123",
  "name": "bucket-2",
  "volume": {
    "name": "fg_oss_1558514455",
    "uuid": "51276f5f-7c6d-11e9-97e8-0050568ea123"
  }
}

```

```

},
"size": 209715200,
"logical_used_size": 157286400,
"encryption": {
  "enabled": false
},
"comment": "S3 bucket.",
"qos_policy": {
  "min_throughput_iops": 0,
  "max_throughput_iops": 1000,
  "max_throughput_mbps": 0,
  "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
  "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
}
}
}
-----

```

=== Creating an S3 bucket for an SVM

The API:

/api/protocols/s3/services/{svm.uuid}/buckets

The call:

```

curl -iku admin:netapp1! -X POST "https://<mgmt-
ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"aggregates\": [ { \"name\":
\"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-8c06-0050568ea123\" } ],
\"comment\": \"S3 bucket.\", \"constituents_per_aggregate\": 4, \"name\":
\"bucket-3\"}"

```

The response:

```

HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
  "num_records": 1,
  "records": [

```



```

{
  "name": "bucket-3",
  "comment": "S3 bucket."
}
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
}
}
-----

```

=== Creating an S3 bucket along with QoS policies for an SVM

The API:

/api/protocols/s3/services/{svm.uuid}/buckets

The call:

```

curl -iku admin:netapp1! -X POST "https://<mgmt-
ip>/api/protocols/s3/services/3e538980-f0af-11e9-ba68-
0050568e9798/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"comment\": \"S3 bucket.\",
\"name\": \"bucket-3\", \"qos_policy\": { \"min_throughput_iops\": 0,
\"max_throughput_iops\": 1000000, \"max_throughput_mbps\": 900000,
\"uuid\": \"02d07a93-6177-11ea-b241-000c293feac8\", \"name\":
\"vs0_auto_gen_policy_02cfa02a_6177_11ea_b241_000c293feac8\" } }"

```

The response:

```

HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-
0050568ea123/buckets/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
  "num_records": 1,
  "records": [

```

```

{
  "name": "bucket-3",
  "comment": "S3 bucket."
}
],
"job": {
  "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
    }
  }
}
}
}
----

```

=== Creating an S3 bucket along with policies for an SVM

The API:

/api/protocols/s3/services/{svm.uuid}/buckets

The call:

```

curl -iku admin:netapp1! -X POST "https://<mgmt-
ip>/api/protocols/s3/services/3e538980-f0af-11e9-ba68-
0050568e9798/buckets?return_records=true" -H "accept: application/json" -H
"Content-Type: application/json" -d "{ \"aggregates\": [ { \"name\":
\"aggr5\", \"uuid\": \"12f3ba4c-7ae0-11e9-8c06-0050568ea123\" } ],
\"comment\": \"S3 bucket.\", \"constituents_per_aggregate\": 4, \"name\":
\"bucket-3\", \"policy\": { \"statements\": [ { \"actions\": [
\"GetObject\" ], \"conditions\": [ { \"operator\": \"ip_address\",
\"source_ips\": [ \"1.1.1.1/23\", \"1.2.2.2/20\" ] }, { \"max_keys\": [
1000 ], \"operator\": \"numeric_equals\" }, { \"delimiters\": [ \"/\" ],
\"operator\": \"string_equals\", \"prefixes\": [ \"pref\" ],
\"usernames\": [ \"user1\" ] } ], \"effect\": \"allow\", \"resources\": [
\"bucket-3/policies/examples/*\" ], \"sid\":
\"AccessToGetObjectForAllUsersofSVM\" }, { \"actions\": [ \"*Object\" ],
\"effect\": \"deny\", \"principals\": [ \"mike\", \"group/grp1\" ],
\"resources\": [ \"bucket-3/policy-docs/*\", \"bucket-3/confidential-*\"
], \"sid\": \"DenyAccessToObjectForMike\" }, { \"actions\": [
\"GetObject\" ], \"effect\": \"allow\", \"principals\": [ \"*\" ],
\"resources\": [ \"bucket-3/readme\" ], \"sid\":
\"AnononymousAccessToGetObjectForUsers\" } ] } }"

```

The response:

```
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:22:14 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/buckets/?name=bucket-3
Content-Length: 353
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "name": "bucket-3",
      "comment": "S3 bucket."
    }
  ],
  "job": {
    "uuid": "2e880171-7e16-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/2e880171-7e16-11e9-bfdc-0050568ea123"
      }
    }
  }
}
```

==== Updating an S3 bucket for an SVM

```
-----

# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/buckets/754389d0-7e13-11e9-bfdc-0050568ea122" -H "accept: application/json?return_records=true" -H "Content-Type: application/json" -d "{ \"comment\": \"Bucket modified.\", \"size\": 111111111111, \"qos_policy\": { \"min_throughput_iops\": 0, \"max_throughput_iops\": 1000000, \"max_throughput_mbps\": 900000, \"uuid\": \"02d07a93-6177-11ea-b241-000c293feac8\", \"name\": \"vs0_auto_gen_policy_02cfa02a_6177_11ea_b241_000c293feac8\"}}}"

# The response:
```

```
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
  "job": {
    "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
      }
    }
  }
}
-----
```

=== Updating an S3 bucket policy for an SVM

```
-----

# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}
```

```
# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/3e538980-f0af-11e9-ba68-0050568e9798/buckets/754389d0-7e13-11e9-bfdc-0050568ea122?return_records=true" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"policy\": { \"statements\": [ { \"actions\": [ \"*\"] , \"conditions\": [ { \"operator\": \"ip_address\", \"source_ips\": [ \"1.1.1.1/23\", \"1.2.2.2/20\" ] }, { \"max_keys\": [ 1000 ], \"operator\": \"numeric_equals\" }, { \"delimiters\": [ \"/\" ], \"operator\": \"string_equals\", \"prefixes\": [ \"pref\" ], \"usernames\": [ \"user1\" ] } ], \"effect\": \"allow\", \"resources\": [ \"*\"] , \"sid\": \"fullAccessForAllPrincipalsToBucket\"} ] } }"
```

```
# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
```

```

{
  "job": {
    "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
      }
    }
  }
}
}
}
----

=== Deleting an S3 bucket policy for an SVM

----

# The API:
/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/3e538980-f0af-11e9-ba68-0050568e9798/buckets/754389d0-7e13-11e9-bfdc-0050568ea122?return_records=true" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"policy\": { \"statements\": [ ] } }"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 May 2019 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
  "job": {
    "uuid": "9beafabb-7e17-11e9-bfdc-0050568ea123",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/9beafabb-7e17-11e9-bfdc-0050568ea123"
      }
    }
  }
}
}
}
----

```

=== Deleting an S3 bucket for a specified SVM

The API:

/api/protocols/s3/services/{svm.uuid}/buckets/{uuid}

The call:

```
curl -iku admin:netapp1! -X DELETE "https://<mgmt-  
ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-  
0050568ea123/buckets/754389d0-7e13-11e9-bfdc-  
0050568ea123?return_records=true" -H "accept: application/json"
```

The response:

```
HTTP/1.1 202 Accepted  
Date: Fri, 24 May 2019 11:40:17 GMT  
Server: libzapid-httpd  
X-Content-Type-Options: nosniff  
Cache-Control: no-cache,no-store,must-revalidate  
Content-Length: 189  
Content-Type: application/json  
{  
  "job": {  
    "uuid": "b3af4a54-7e18-11e9-bfdc-0050568ea123",  
    "_links": {  
      "self": {  
        "href": "/api/cluster/jobs/b3af4a54-7e18-11e9-bfdc-0050568ea123"  
      }  
    }  
  }  
}  
}
```

[[ID913a9fbec954c4bd4115f7a255844643]]

= Retrieve S3 bucket configurations for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/buckets`#

Introduced In: 9.7

Retrieves the S3 bucket's configuration of an SVM. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to

'**'.

== Related ONTAP commands

- * ``vserver object-store-server bucket show``
- * ``vserver object-store-server bucket policy statement show``
- * ``vserver object-store-server bucket policy-statement-condition show``

== Learn more

*
xref:{relative_path}protocols_s3_services_svm.uuid_buckets_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/buckets]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|policy.statements.conditions.source_ips

|string

|query

|False

a|Filter by policy.statements.conditions.source_ips

* Introduced in: 9.8

|policy.statements.conditions.prefixes

|string

|query

|False

a|Filter by policy.statements.conditions.prefixes

* Introduced in: 9.8

|policy.statements.conditions.operator

|string

|query

```

|False
a|Filter by policy.statements.conditions.operator

* Introduced in: 9.8


|policy.statements.conditions.max_keys
|integer
|query
|False
a|Filter by policy.statements.conditions.max_keys

* Introduced in: 9.8


|policy.statements.conditions.usernames
|string
|query
|False
a|Filter by policy.statements.conditions.usernames

* Introduced in: 9.8


|policy.statements.conditions.delimiters
|string
|query
|False
a|Filter by policy.statements.conditions.delimiters

* Introduced in: 9.8


|policy.statements.principals
|string
|query
|False
a|Filter by policy.statements.principals

* Introduced in: 9.8


|policy.statements.sid
|string
|query
|False
a|Filter by policy.statements.sid

```


* Introduced in: 9.8

```
|policy.statements.effect
|string
|query
|False
a|Filter by policy.statements.effect
```

* Introduced in: 9.8

```
|policy.statements.actions
|string
|query
|False
a|Filter by policy.statements.actions
```

* Introduced in: 9.8

```
|policy.statements.resources
|string
|query
|False
a|Filter by policy.statements.resources
```

* Introduced in: 9.8

```
|logical_used_size
|integer
|query
|False
a|Filter by logical_used_size
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
```

```

|query
|False
a|Filter by svm.name

|encryption.enabled
|boolean
|query
|False
a|Filter by encryption.enabled

|qos_policy.max_throughput_mbps
|integer
|query
|False
a|Filter by qos_policy.max_throughput_mbps

* Introduced in: 9.8

|qos_policy.max_throughput_iops
|integer
|query
|False
a|Filter by qos_policy.max_throughput_iops

* Introduced in: 9.8

|qos_policy.uuid
|string
|query
|False
a|Filter by qos_policy.uuid

* Introduced in: 9.8

|qos_policy.min_throughput_iops
|integer
|query
|False
a|Filter by qos_policy.min_throughput_iops

* Introduced in: 9.8

```

```
|qos_policy.min_throughput_mbps
|integer
|query
|False
a|Filter by qos_policy.min_throughput_mbps
```

* Introduced in: 9.8

```
|qos_policy.name
|string
|query
|False
a|Filter by qos_policy.name
```

* Introduced in: 9.8

```
|volume.uuid
|string
|query
|False
a|Filter by volume.uuid
```

```
|volume.name
|string
|query
|False
a|Filter by volume.name
```

```
|size
|integer
|query
|False
a|Filter by size
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
```

```

|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#s3_bucket_svm[s3_bucket_svm]]
a|

|===

```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "comment": "S3 bucket.",
      "logical_used_size": 0,
      "name": "bucket1",
      "policy": {
        "statements": [
          {
            "actions": [
              "GetObject",
              "PutObject",
              "DeleteObject",
              "ListBucket"
            ],
            "conditions": [
              {
                "delimiters": [
                  "/"
                ],
                "max_keys": [
                  "1000"
                ],
                "operator": "ip_address",
                "prefixes": [
                  "pref"
                ],
                "source_ips": [
                  "1.1.1.1",
                  "1.2.2.0/24"
                ],
                "usernames": [
                  "user1"
                ]
              }
            ]
          }
        ]
      }
    }
  ]
}
```

```

        ]
    }
},
"effect": "allow",
"principals": [
    "user1",
    "group/grp1"
],
"resources": [
    "bucket1",
    "bucket1/*"
],
"sid": "FullAccessToUser1"
}
]
},
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    }
},

```

```

        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

```



```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#encryption]
[.api-collapsible-fifth-title]
encryption

[cols=3*,options=header]
|===

```

Name	Type	Description
enabled	boolean	a Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket.

===
[#s3_bucket_policy_condition] [.api-collapsible-fifth-title] s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

[cols=3*,options=header]	===	Name	Type	Description
delimiters	array[string]	a An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.		
max_keys	array[integer]	a An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.		
operator	string	a Condition operator that is applied to the specified condition key.		
prefixes	array[string]			

a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips

|array[string]

a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames

|array[string]

a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]

[.api-collapsible-fifth-title]

s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|actions

|array[string]

a|

|conditions

|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]

a|Specifies bucket policy conditions.

|effect

|string

a|Specifies whether access is allowed or denied when a user requests the specific action. If access (to allow) is not granted explicitly to a

resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.

```
|principals
|array[string]
a|
```

```
|resources
|array[string]
a|
```

```
|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.
```

```
|===
```

```
[#policy]
[.api-collapsible-fifth-title]
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

```
|===
```

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|max_throughput_iops
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
```

```
|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
```

```
|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
```

```
|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
```

This is mutually exclusive with name and UUID during POST and PATCH.

```
|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.
```

```
|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.
```

|===

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

|===

```
[#volume]
[.api-collapsible-fifth-title]
```

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket_svm]
```

```
[.api-collapsible-fifth-title]
```

```
s3_bucket_svm
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```

|comment
|string
a|Can contain any additional information about the bucket being created or
modified.

|encryption
|link:#encryption[encryption]
a|

|logical_used_size
|integer
a|Specifies the bucket logical used size up to this point.

|name
|string
a|Specifies the name of the bucket. Bucket name is a string that can only
contain the following combination of ASCII-range alphanumeric characters
0-9, a-z, ".", and "-".

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifies "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifies "min_throughput_iops" is only supported on volumes hosted on a
node that is flash optimized. A pre-created QoS policy can also be used by
specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or
assigning a QoS policy to a bucket is not supported if its containing
volume or SVM already has a QoS policy attached.

|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|Specifies the unique identifier of the bucket.
```

```
|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====

[[IDe68227d0a504aa7119033a61d9168cf4]]
= Create an S3 bucket configuration for an SVM

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/buckets`#

*Introduced In:* 9.7

Creates the S3 bucket configuration of an SVM.

== Important notes

* Each SVM can have one or more bucket configurations.
* Aggregate lists should be specified explicitly. If not specified, then

```

the bucket is auto-provisioned as a FlexGroup.

- * Constituents per aggregate specifies the number of components (or FlexVols) per aggregate. Is specified only when an aggregate list is explicitly defined.

- * An access policy can be created when a bucket is created.

- * "qos_policy" can be specified if a bucket needs to be attached to a QoS group policy during creation time.

== Required properties

- * `svm.uuid` - Existing SVM in which to create the bucket configuration.

- * `name` - Bucket name that is to be created.

== Recommended optional properties

- * `aggregates` - List of aggregates for the FlexGroup on which the bucket is hosted on.

- * `constituents_per_aggregate` - Number of constituents per aggregate.

- * `size` - Specifying the bucket size is recommended.

- * `policy` - Specifying policy enables users to perform operations on buckets. Hence specifying the resource permissions is recommended.

- * `qos_policy` - A QoS policy for buckets.

== Default property values

- * `size` - 800MB

- * `comment` - ""

- * `aggregates` - No default value.

- * `constituents_per_aggregate` - `_4_` , if an aggregates list is specified. Otherwise, no default value.

- * `policy.statements.actions` - GetObject, PutObject, DeleteObject, ListBucket.

- * `policy.statements.principals` - all S3 users and groups in the SVM.

- * `policy.statements.resources` - all objects in the bucket.

- * `policy.statements.conditions` - list of bucket policy conditions.

- * `qos-policy` - No default value.

== Related ONTAP commands

- * `vserver object-store-server bucket create`

- * `vserver object-store-server bucket policy statement create`

- * `vserver object-store-server bucket policy-statement-condition create`

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_buckets_endpoint_overvi

```
ew.html[DOC /protocols/s3/services/{svm.uuid}/buckets]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|return_timeout
```

```
|integer
```

```
|query
```

```
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

```
|return_records
```

```
|boolean
```

```
|query
```

```
|False
```

a|The default is false. If set to true, the records are returned.

* Default value:

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

a|UUID of the SVM to which this object belongs.

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|aggregates
```

```
|array[link:#aggregates[aggregates]]
```

a|A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume.

```
|comment
```

```
|string
```

a|Can contain any additional information about the bucket being created or modified.

```
|constituents_per_aggregate
```

```
|integer
```

a|Specifies the number of constituents or FlexVol volumes per aggregate. A FlexGroup volume consisting of all such constituents across all specified aggregates is created. This option is used along with the aggregates option and cannot be used independently.

```
|encryption
```

```
|link:#encryption[encryption]
```

```
a|
```

```
|logical_used_size
```

```
|integer
```

a|Specifies the bucket logical used size up to this point.

```
|name
```

```
|string
```

a|Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-".

```
|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.
```

```
|qos_policy
|link:#qos_policy[qos_policy]
a|Specifies "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifies "min_throughput_iops" is only supported on volumes hosted on a
node that is flash optimized. A pre-created QoS policy can also be used by
specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or
assigning a QoS policy to a bucket is not supported if its containing
volume or SVM already has a QoS policy attached.
```

```
|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.
```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|Specifies the unique identifier of the bucket.
```

```
|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.
```

```
|===
```

```
.Example request
[%collapsible%closed]
=====
```

```
[source,json,subs=+macros]
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "comment": "S3 bucket.",
  "constituents_per_aggregate": "4",
  "logical_used_size": 0,
  "name": "bucket1",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              "1000"
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
            ]
          }
        ],
        "effect": "allow",
        "principals": [
          "user1",
          "group/grp1"
        ]
      }
    ]
  }
}
```



```

    ],
    "resources": [
        "bucket1",
        "bucket1/*"
    ],
    "sid": "FullAccessToUser1"
}
]
},
"qos_policy": {
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
====

== Response

```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
=====

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error code | Message
//end header
//end row
//start row
|92405777 +
//end row
//start row
|"Failed to create bucket "{bucket name}" for SVM "{svm.name}". Reason:
{Reason of failure}. ";
//end row
//start row
|92405785 +

```

```

//end row
//start row
|"Bucket name "{bucket name}" contains invalid characters. Valid
characters for a bucket name are 0-9, a-z, ".", and "-". ";
//end row
//start row
|92405786 +
//end row
//start row
|"Bucket name "{bucket name}" is not valid. Bucket names must have between
3 and 63 characters. ";
//end row
//start row
|92405811 +
//end row
//start row
|"Failed to create bucket "{bucket name}" for SVM "{svm.name}". Wait a few
minutes and try the operation again.";
//end row
//start row
|92405812 +
//end row
//start row
|"Failed to create the object store volume. Reason: {Reason for
failure}.";
//end row
//start row
|92405819 +
//end row
//start row
|"Cannot provision an object store server volume for bucket "{bucket
name}" in SVM "{svm.name}" on the following aggregates because they are
SnapLock aggregates: {List of aggregates.name}.";
//end row
//start row
|92405820 +
//end row
//start row
|"Failed to check whether the aggregate "{aggregates.name}" is a
FabricPool. Reason: {Reason for failure}.";
//end row
//start row
|92405821 +
//end row
//start row
|"Cannot provision an object store server volume for bucket "{bucket

```

```

name}" in SVM "{svm.name}" on the following aggregates because they are
FabricPool: {List of aggregates.name}.";
//end row
//start row
|92405827 +
//end row
//start row
|"Internal Error. Unable to generate object store volume name.";
//end row
//start row
|92405857 +
//end row
//start row
|"One or more aggregates must be specified if "constituents_per_aggregate"
is specified.";
//end row
//start row
|92405858 +
//end row
//start row
|"Failed to "create" the "bucket" because the operation is only supported
on data SVMs.";
//end row
//start row
|92405859 +
//end row
//start row
|"The specified "aggregates.uuid" "{aggregates.uuid}" does not exist.";
//end row
//start row
|92405860 +
//end row
//start row
|"The specified "aggregates.name" "{aggregates.name}" and
"aggregates.uuid" "{aggregates.uuid}" refer to different aggregates.";
//end row
//start row
|92405861 +
//end row
//start row
|"The specified SVM UUID or bucket UUID does not exist.";
//end row
//start row
|92405863 +
//end row
//start row

```

```

|"An error occurs when creating an access policy. The reason for failure
is detailed in the error message.";
//end row
//start row
|92405891 +
//end row
//start row
|The resources specified in the access policy are not valid. Valid ways to
specify a resource are *, <bucket-name>, <bucket-name>/.../.... Valid
characters for a resource are 0-9, A-Z, a-z, _, +, comma, ;, :, =, ., &,
@,?, (, ), single quote, *, !, - and $.

//end row
//start row
|92405894 +
//end row
//start row
|"Statements, principals and resources list can have a maximum of 10
entries.";
//end row
//start row
|92405897 +
//end row
//start row
|The principals specified in the access policy are not in the correct
format. User name must be in between 1 and 64 characters. Valid characters
for a user name are 0-9, A-Z, a-z, _, +, =, comma, ., @, and - .

//end row
//start row
|92405898 +
//end row
//start row
|"The SID specified in the access policy is not valid. Valid characters
for a SID are 0-9, A-Z and a-z.";
//end row
|===
//end table

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#encryption]
[.api-collapsible-fifth-title]

```

encryption

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket.

|===

[#s3_bucket_policy_condition]

[.api-collapsible-fifth-title]

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|delimiters

|array[string]

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

|max_keys

|array[integer]

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

|operator

|string

a|Condition operator that is applied to the specified condition key.

|prefixes
|array[string]
a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips
|array[string]
a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames
|array[string]
a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]
[.api-collapsible-fifth-title]
s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|actions
|array[string]
a|

|conditions
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
a|Specifies bucket policy conditions.


```
|effect
|string
a|Specifies whether access is allowed or denied when a user requests the
specific action. If access (to allow) is not granted explicitly to a
resource, access is implicitly denied. Access can also be denied
explicitly to a resource, in order to make sure that a user cannot access
it, even if a different policy grants access.
```

```
|principals
|array[string]
a|
```

```
|resources
|array[string]
a|
```

```
|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.
```

```
|===
```

```
[#policy]
[.api-collapsible-fifth-title]
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

|===

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|max_throughput_iops
|integer
```

a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|max_throughput_mbps
|integer
```

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_iops
|integer
```

a|Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_mbps
|integer
```

a|Specifies the minimum throughput in Megabytes per sec, 0 means none.

This is mutually exclusive with name and UUID during POST and PATCH.

```
|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.
```

```
|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.
```

|===

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

|===

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket_svm]
```

```
[.api-collapsible-fifth-title]
```

```
s3_bucket_svm
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|aggregates
```

```
|array[link:#aggregates[aggregates]]
```

```
a|A list of aggregates for FlexGroup volume constituents where the bucket is hosted. If this option is not specified, the bucket is auto-provisioned as a FlexGroup volume.
```

```

|comment
|string
a|Can contain any additional information about the bucket being created or
modified.

|constituents_per_aggregate
|integer
a|Specifies the number of constituents or FlexVol volumes per aggregate. A
FlexGroup volume consisting of all such constituents across all specified
aggregates is created. This option is used along with the aggregates
option and cannot be used independently.

|encryption
|link:#encryption[encryption]
a|

|logical_used_size
|integer
a|Specifies the bucket logical used size up to this point.

|name
|string
a|Specifies the name of the bucket. Bucket name is a string that can only
contain the following combination of ASCII-range alphanumeric characters
0-9, a-z, ".", and "-".

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifes "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifes "min_throughput_iops" is only supported on volumes hosted on a
node that is flash optimized. A pre-created QoS policy can also be used by
specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or

```

assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

|size

|integer

a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|Specifies the unique identifier of the bucket.

|volume

|link:#volume[volume]

a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

|===

[#job_link]

[.api-collapsible-fifth-title]

job_link

[cols=3*,options=header]

|===

|Name

|Type

|Description

|uuid

|string

a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID2af19b2436466f1fab396b1e15ce0841]]
= Delete the S3 bucket configuration for an SVM

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/buckets/{uuid}`#

*Introduced In:* 9.7

Deletes the S3 bucket configuration of an SVM. An access policy is also
deleted on an S3 bucket "delete" command.

== Related ONTAP commands

* `vserver object-store-server bucket delete`
* `vserver object-store-server bucket policy statement delete`
* `vserver object-store-server bucket policy-statement-condition delete`

== Learn more

*
xref:{relative_path}protocols_s3_services_svm.uuid_buckets_endpoint_overvi
ew.html[DOC /protocols/s3/services/{svm.uuid}/buckets]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

```



```
|uuid
|string
|path
|True
a|The unique identifier of the bucket.
```

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds. This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job. If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.
```

```
* Default value: 1
* Max value: 120
* Min value: 0
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|===
```

```
== Response
```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
=====

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error code | Message
//end header
//end row
//start row
|92405811 +
//end row
//start row

```

```

|"Failed to delete bucket "{bucket name}" for SVM "{svm.name}". Wait a few
minutes and try the operation again.";
//end row
//start row
|92405858 +
//end row
//start row
|"Failed to "delete" the "bucket" because the operation is only supported
on data SVMs.";
//end row
//start row
|92405861 +
//end row
//start row
|"The specified SVM UUID or bucket UUID does not exist.";
//end row
//start row
|92405779 +
//end row
//start row
|"Failed to remove bucket "{bucket name}" for SVM "{svm.name}". Reason:
{Reason for failure}.";
//end row
//start row
|92405813 +
//end row
//start row
|"Failed to delete the object store volume. Reason: {Reason for
failure}.";
//end row
//start row
|92405864 +
//end row
//start row
|"An error occurred when deleting an access policy. The reason for failure
is detailed in the error message.";
//end row
|===
//end table

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href

```

```

|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

[#error_arguments]

```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID7b8041ba82feaefc42b067f71239e066]]
= Retrieve an S3 bucket for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/buckets/{uuid}`#
```

Introduced In: 9.7

Retrieves the S3 bucket configuration of an SVM. Note that in order to retrieve S3 bucket policy conditions, the 'fields' option should be set to '***'.

== Related ONTAP commands

```
* `vserver object-store-server bucket show`
* `vserver object-store-server bucket policy statement show`
* `vserver object-store-server bucket policy-statement-condition show`
```

== Learn more

```
*
xref:{relative_path}protocols_s3_services_svm.uuid_buckets_endpoint_overvi
ew.html[DOC /protocols/s3/services/{svm.uuid}/buckets]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
```

```

|Description

|uuid
|string
|path
|True
a|The unique identifier of the bucket.

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Can contain any additional information about the bucket being created or
modified.

|encryption
|link:#encryption[encryption]
a|

|logical_used_size
|integer

```


a|Specifies the bucket logical used size up to this point.

|name

|string

a|Specifies the name of the bucket. Bucket name is a string that can only contain the following combination of ASCII-range alphanumeric characters 0-9, a-z, ".", and "-".

|policy

|link:#policy[policy]

a|A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

|qos_policy

|link:#qos_policy[qos_policy]

a|Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

|size

|integer

a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|Specifies the unique identifier of the bucket.

|volume

|link:#volume[volume]

a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

|===

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "comment": "S3 bucket.",
  "logical_used_size": 0,
  "name": "bucket1",
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
              "/"
            ],
            "max_keys": [
              "1000"
            ],
            "operator": "ip_address",
            "prefixes": [
              "pref"
            ],
            "source_ips": [
              "1.1.1.1",
              "1.2.2.0/24"
            ],
            "usernames": [
              "user1"
            ]
          }
        ],
        "effect": "allow",
        "principals": [
```

```

        "user1",
        "group/grp1"
    ],
    "resources": [
        "bucket1",
        "bucket1/*"
    ],
    "sid": "FullAccessToUser1"
}
]
},
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
=====

```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
```

```

href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]

```

```

a|

|name
|string
a|

|uuid
|string
a|

|===

[#encryption]
[.api-collapsible-fifth-title]
encryption

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Specifies whether encryption is enabled on the bucket. By default,
encryption is disabled on a bucket.

|===

[#s3_bucket_policy_condition]
[.api-collapsible-fifth-title]
s3_bucket_policy_condition

Information about policy conditions based on various condition operators
and condition keys.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|delimiters

```

|array[string]

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

|max_keys

|array[integer]

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

|operator

|string

a|Condition operator that is applied to the specified condition key.

|prefixes

|array[string]

a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips

|array[string]

a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames

|array[string]

a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]

[.api-collapsible-fifth-title]

s3_bucket_policy_statement

Specifies information about a single access permission.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|actions
|array[string]
a|

|conditions
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
a|Specifies bucket policy conditions.

|effect
|string
a|Specifies whether access is allowed or denied when a user requests the
specific action. If access (to allow) is not granted explicitly to a
resource, access is implicitly denied. Access can also be denied
explicitly to a resource, in order to make sure that a user cannot access
it, even if a different policy grants access.

|principals
|array[string]
a|

|resources
|array[string]
a|

|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.

|===

[#policy]
[.api-collapsible-fifth-title]
policy

```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when

an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```

```
|===
```

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|max_throughput_iops
|integer
a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
```

|max_throughput_mbps
|integer
a|Specifies the maximum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.

|min_throughput_iops
|integer
a|Specifies the minimum throughput in IOPS, 0 means none. Setting
"min_throughput" is supported on AFF platforms only, unless FabricPool
tiering policies are set. This is mutually exclusive with name and UUID
during POST and PATCH.

|min_throughput_mbps
|integer
a|Specifies the minimum throughput in Megabytes per sec, 0 means none.
This is mutually exclusive with name and UUID during POST and PATCH.

|name
|string
a|The QoS policy group name. This is mutually exclusive with UUID and
other QoS attributes during POST and PATCH.

|uuid
|string
a|The QoS policy group UUID. This is mutually exclusive with name and
other QoS attributes during POST and PATCH.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume.
```

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.
```

* example: 028baa66-41bd-11e9-81d5-00a0986138f7

* Introduced in: 9.6

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID226a367df40b86bb6e09489d4dca37bf]]
= Update an S3 bucket configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/buckets/{uuid}`#
```

```
*Introduced In:* 9.7
```

Updates the S3 bucket configuration of an SVM.

== Important notes

* The following fields can be modified for a bucket:

*** `comment` - Any information related to the bucket.

*** `size` - Bucket size.

*** `policy` - An access policy for resources (buckets and objects) that defines their permissions. New policies are created after existing policies are deleted. To retain any of the existing policy statements, you need to specify those statements again. Policy conditions can also be modified using this API.

*** `qos_policy` - A QoS policy for buckets.

== Related ONTAP commands

```
* `vserver object-store-server bucket modify`  
* `vserver object-store-server bucket policy statement modify`  
* `vserver object-store-server bucket policy-statement-condition modify`
```

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_buckets_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/buckets]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|The unique identifier of the bucket.

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

|svm.uuid

```

|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Can contain any additional information about the bucket being created or
modified.

|encryption
|link:#encryption[encryption]
a|

|logical_used_size
|integer
a|Specifies the bucket logical used size up to this point.

|policy
|link:#policy[policy]
a|A policy is an object associated with a bucket. It defines resource
(bucket, folder, or object) permissions. These policies get evaluated when
an S3 user makes a request by executing a specific command. The user must
be part of the principal (user or group) specified in the policy.
Permissions in the policies determine whether the request is allowed or
denied.

|qos_policy
|link:#qos_policy[qos_policy]
a|Specifes "qos_policy.max_throughput_iops" and/or
"qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".
Specifes "min_throughput_iops" is only supported on volumes hosted on a

```

node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
|size
|integer
a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.
```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|Specifies the unique identifier of the bucket.
```

```
|volume
|link:#volume[volume]
a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.
```

```
|===
```

```
.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "comment": "S3 bucket.",
  "logical_used_size": 0,
  "policy": {
    "statements": [
      {
        "actions": [
          "GetObject",
          "PutObject",
          "DeleteObject",
          "ListBucket"
        ],
        "conditions": [
          {
            "delimiters": [
```



```

        "/"
    ],
    "max_keys": [
        "1000"
    ],
    "operator": "ip_address",
    "prefixes": [
        "pref"
    ],
    "source_ips": [
        "1.1.1.1",
        "1.2.2.0/24"
    ],
    "usernames": [
        "user1"
    ]
}
],
"effect": "allow",
"principals": [
    "user1",
    "group/grp1"
],
"resources": [
    "bucket1",
    "bucket1/*"
],
"sid": "FullAccessToUser1"
}
]
},
"qos_policy": {
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "1677721600",
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "414b29a1-3b26-11e9-bd58-0050568ea055",
"volume": {

```

```

    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
====

== Response

```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error code | Message
//end header

```

```

//end row
//start row
|92405778 +
//end row
//start row
|"Failed to modify bucket "{bucket name}" for SVM "{svm.name}". Reason:
{Reason for failure}. ";
//end row
//start row
|92405846 +
//end row
//start row
|"Failed to modify the object store volume. Reason: {Reason for
failure}. ";
//end row
//start row
|92405811 +
//end row
//start row
|"Failed to modify bucket "{bucket name}" for SVM "{svm.name}". Wait a few
minutes and try the operation again.";
//end row
//start row
|92405858 +
//end row
//start row
|"Failed to "modify" the "bucket" because the operation is only supported
on data SVMs.";
//end row
//start row
|92405861 +
//end row
//start row
|"The specified SVM UUID or bucket UUID does not exist.";
//end row
//start row
|92405863 +
//end row
//start row
|"An error occurs when creating an access policy. The reason for failure
is detailed in the error message.";
//end row
//start row
|92405864 +
//end row
//start row

```

```

|"An error occurs when deleting an access policy. The reason for failure
is detailed in the error message.";
//end row
//start row
|92405891 +
//end row
//start row
|The resources specified in the access policy are not valid. Valid ways to
specify a resource are *, <bucket-name>, <bucket-name>/.../.... Valid
characters for a resource are 0-9, A-Z, a-z, _, +, comma, ;, :, =, ., &,
@,?, (, ), single quote, *, !, - and $.

//end row
//start row
|92405894 +
//end row
//start row
|"Statements, principals and resources list can have a maximum of 10
entries.";
//end row
//start row
|92405897 +
//end row
//start row
|The principals specified in the access policy are not in the correct
format. User name must be in between 1 and 64 characters. Valid characters
for a user name are 0-9, A-Z, a-z, _, +, =, comma, ., @, and - .

//end row
//start row
|92405898 +
//end row
//start row
|"The SID specified in the access policy is not valid. Valid characters
for a SID are 0-9, A-Z and a-z.";
//end row
|===
//end table

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#encryption]
[.api-collapsible-fifth-title]

```

encryption

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|Specifies whether encryption is enabled on the bucket. By default, encryption is disabled on a bucket.

|===

[#s3_bucket_policy_condition]

[.api-collapsible-fifth-title]

s3_bucket_policy_condition

Information about policy conditions based on various condition operators and condition keys.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|delimiters

|array[string]

a|An array of delimiters that are compared with the delimiter value specified at the time of execution of an S3-based command, using the condition operator specified.

|max_keys

|array[integer]

a|An array of maximum keys that are allowed or denied to be retrieved using an S3 list operation, based on the condition operator specified.

|operator

|string

a|Condition operator that is applied to the specified condition key.

|prefixes
|array[string]
a|An array of prefixes that are compared with the input prefix value specified at the time of execution of an S3-based command, using the condition operator specified.

|source_ips
|array[string]
a|An array of IP address ranges that are compared with the IP address of a source command at the time of execution of an S3-based command, using the condition operator specified.

|usernames
|array[string]
a|An array of usernames that a current user in the context is evaluated against using the condition operators.

|===

[#s3_bucket_policy_statement]
[.api-collapsible-fifth-title]
s3_bucket_policy_statement

Specifies information about a single access permission.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|actions
|array[string]
a|

|conditions
|array[link:#s3_bucket_policy_condition[s3_bucket_policy_condition]]
a|Specifies bucket policy conditions.

```
|effect
|string
a|Specifies whether access is allowed or denied when a user requests the
specific action. If access (to allow) is not granted explicitly to a
resource, access is implicitly denied. Access can also be denied
explicitly to a resource, in order to make sure that a user cannot access
it, even if a different policy grants access.
```

```
|principals
|array[string]
a|
```

```
|resources
|array[string]
a|
```

```
|sid
|string
a|Specifies the statement identifier used to differentiate between
statements.
```

```
|===
```

```
[#policy]
[.api-collapsible-fifth-title]
policy
```

A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|statements
|array[link:#s3_bucket_policy_statement[s3_bucket_policy_statement]]
a|Specifies bucket access policy statement.
```



```
|===
```

```
[#qos_policy]  
[.api-collapsible-fifth-title]  
qos_policy
```

Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|max_throughput_iops  
|integer
```

a|Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|max_throughput_mbps  
|integer
```

a|Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_iops  
|integer
```

a|Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

```
|min_throughput_mbps  
|integer
```

a|Specifies the minimum throughput in Megabytes per sec, 0 means none.

This is mutually exclusive with name and UUID during POST and PATCH.

|name

|string

a|The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.

|uuid

|string

a|The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#volume]

[.api-collapsible-fifth-title]

volume

Specifies the FlexGroup volume name and UUID where the bucket is hosted.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the volume.
```

```
|uuid
```

```
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#s3_bucket_svm]
```

```
[.api-collapsible-fifth-title]
```

```
s3_bucket_svm
```

A bucket is a container of objects. Each bucket defines an object namespace. S3 requests specify objects using a bucket-name and object-name pair. An object resides within a bucket.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|comment
```

```
|string
```

```
a|Can contain any additional information about the bucket being created or modified.
```

```
|encryption
```

|link:#encryption[encryption]

a|

|logical_used_size

|integer

a|Specifies the bucket logical used size up to this point.

|policy

|link:#policy[policy]

a|A policy is an object associated with a bucket. It defines resource (bucket, folder, or object) permissions. These policies get evaluated when an S3 user makes a request by executing a specific command. The user must be part of the principal (user or group) specified in the policy. Permissions in the policies determine whether the request is allowed or denied.

|qos_policy

|link:#qos_policy[qos_policy]

a|Specifies "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops". Specifies "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties. Setting or assigning a QoS policy to a bucket is not supported if its containing volume or SVM already has a QoS policy attached.

|size

|integer

a|Specifies the bucket size in bytes; ranges from 80MB to 64TB.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|Specifies the unique identifier of the bucket.

|volume

|link:#volume[volume]

a|Specifies the FlexGroup volume name and UUID where the bucket is hosted.

|===

[#job_link]

[.api-collapsible-fifth-title]

job_link

[cols=3*,options=header]

|===

|Name

|Type

|Description

|uuid

|string

a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

```

[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage S3 service group configuration

:leveloffset: +1

[[IDbbc197f9a3c05d41dacad18bd39dac3f]]
= Protocols S3 services svm.uuid groups endpoint overview

```

== Overview

An S3 group consists of one or many users. Policies are attached to the S3 group to have access control over S3 resources at group level.

== Examples

=== Retrieving all fields for all S3 groups of an SVM

The API:

/api/protocols/s3/services/{svm.uuid}/groups:

The call:

```
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/groups?fields=*&return_records=true&return_timeout=15" -H "accept: application/json"
```

The response:

```
{
  "records": [
    {
      "comment": "Admin group",
      "id": 5,
      "name": "Admin-Group",
      "policies": [
        { "name": "Policy1" },
        { "name": "Policy2" },
        { "name": "Policy3" }
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "users": [
        { "name": "User1" },
        { "name": "User2" },
        { "name": "User3" }
      ]
    },
    {

```

```

    "comment": "Admin group",
    "id": 6,
    "name": "Admin-Group1",
    "policies": [
      { "name": "Policy1" },
      { "name": "Policy2" },
      { "name": "Policy3" }
    ],
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "users": [
      { "name": "User1" },
      { "name": "User2" },
      { "name": "User6" }
    ]
  }
],
"num_records": 2
}
-----

```

=== Retrieving the specified group in the SVM

```

-----

# The API:
/api/protocols/s3/services/{svm.uuid}/groups/{group_id}:

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/groups/5?fields=*" -H "accept: application/json"

# The response:
{
  "comment": "Admin group",
  "id": 5,
  "name": "Admin-Group",
  "policies": [
    { "name": "Policy1" },
    { "name": "Policy2" },
    { "name": "Policy3" }
  ],
  "svm": {
    "name": "svm1",

```



```

    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    { "name": "User1" },
    { "name": "User2" },
    { "name": "User3" }
  ]
}
----

=== Creating an S3 group for an SVM

----

# The API:
/api/protocols/s3/services/{svm.uuid}/groups

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/groups?return_records=true" -H "accept: application/json" -H "authorization: Basic YWRtaW46bmV0YXBwMSE=" -H "Content-Type: application/json" -d "{ \"comment\": \"Admin group\", \"name\": \"Admin-Group\", \"policies\": [ { \"name\": \"Policy1\" }, { \"name\": \"Policy2\" }, { \"name\": \"Policy3\" } ], \"users\": [ { \"name\": \"User1\" }, { \"name\": \"User2\" }, { \"name\": \"User3\" } ]}"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 289
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "comment": "Admin group",
      "id": 5,
      "name": "Admin-Group",
      "policies": [
        { "name": "Policy1" },
        { "name": "Policy2" },
        { "name": "Policy3" }
      ]
    }
  ]
}

```

```

    ],
    "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "users": [
        { "name": "User1" },
        { "name": "User2" },
        { "name": "User3" }
    ]
}
]
}
}
----

=== Updating an S3 group for an SVM

----

# The API:
/api/protocols/s3/services/{svm.uuid}/groups/{group_id}:

# The call:
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/groups/5" -H "accept: application/json" -H
"authorization: Basic YWRtaW46bmV0YXBwMSE=" -H "Content-Type:
application/json" -d "{ \"comment\": \"Admin group\", \"name\": \"Admin-
Group\", \"policies\": [ { \"name\": \"Policy1\" } ], \"users\": [ {
\"name\": \"user-1\" } ]}"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
----

=== Deleting an S3 group for a specified SVM

----

# The API:
/api/protocols/s3/services/{svm.uuid}/groups/{group_id}:

```

```

# The call:
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/groups/5" -H "accept: application/json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
-----

[[ID3a1c693164313f04689f9f424c81d97a]]
= Retrieve an S3 group configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/groups`#

*Introduced In:* 9.8

Retrieves the S3 group's SVM configuration.

== Related ONTAP commands

* `vserver object-store-server group show`

== Learn more

*
xref:{relative_path}protocols_s3_services_svm.uuid_groups_endpoint_overvie
w.html[DOC /protocols/s3/services/{svm.uuid}/groups]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In

```

```
|Required
|Description

|comment
|string
|query
|False
a|Filter by comment

|name
|string
|query
|False
a|Filter by name

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|id
|integer
|query
|False
a|Filter by id

|users.name
|string
|query
|False
a|Filter by users.name

|policies.name
|string
```

```
|query
|False
a|Filter by policies.name
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|
```

```
|num_records
|integer
a|Number of records
```

```
|records
|array[link:#s3_group[s3_group]]
a|
```

|===

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
```

```

    },
    "self": {
        "href": "/api/resourcelink"
    }
},
"records": [
    {
        "comment": "Admin group",
        "id": "5",
        "name": "Admin-Group",
        "policies": [
            {
                "_links": {
                    "self": {
                        "href": "/api/resourcelink"
                    }
                },
                "name": "Policy1"
            }
        ],
        "svm": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "users": [
            {
                "_links": {
                    "self": {
                        "href": "/api/resourcelink"
                    }
                },
                "name": "user-1"
            }
        ]
    }
]
}
====

== Error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]

```



```

|===
|Name
|Type
|Description

|href
|string
a|

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

```

```

|===

[#policies]
[.api-collapsible-fifth-title]
policies

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Specifies the name of the policy. A policy name length can range from 1
to 128 characters and can only contain the following combination of
characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid

```

```
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#users]
[.api-collapsible-fifth-title]
users
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|Specifies the name of the user. A user name length can range from 1 to
64 characters and can only contain the following combination of characters
0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".
```

```
|===
```

```
[#s3_group]
[.api-collapsible-fifth-title]
s3_group
```

This is a container for S3 user groups.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|comment
|string
a|Can contain any additional information about the group being created or
```

modified.

|id

|integer

a|Specifies a unique group ID used to identify a particular group. This parameter should not be specified in the POST method. A group ID is automatically generated and it is retrieved using the GET method. Group id is SVM scoped.

|name

|string

a|Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|policies

|array[link:#policies[policies]]

a|Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies.

|svm

|link:#svm[svm]

a|

|users

|array[link:#users[users]]

a|Specifies the list of users who belong to the group.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

```

|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

[[ID484caacebb2728fb2041c981fe237c82]]

= Create an S3 group configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/protocols/s3/services/{svm.uuid}/groups`#

Introduced In: 9.8

Creates the S3 group configuration.

== Important notes

* Each SVM can have one or more s3 group configurations.

== Required properties

* `svm.uuid` - Existing SVM in which to create the user configuration.

* `name` - Group name that is to be created.

* `users` - List of users to be added into the group.

* `policies` - List of policies are to be attached to this group.

== Recommended optional properties

* `comment` - Short description about the S3 Group.

== Related ONTAP commands

* `vserver object-store-server group create`

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_groups_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/groups]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

```
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.
```

* Default value:

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|comment
|string
a|Can contain any additional information about the group being created or modified.
```

```
|id
|integer
a|Specifies a unique group ID used to identify a particular group. This parameter should not be specified in the POST method. A group ID is automatically generated and it is retrieved using the GET method. Group id is SVM scoped.
```

```
|name
|string
a|Specifies the name of the group. A group name length can range from 1 to
```

128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

```
|policies
|array[link:#policies[policies]]
a|Specifies a list of policies that are attached to the group. The
wildcard character "*" is a valid value for specifying all policies.
```

```
|svm
|link:#svm[svm]
a|
```

```
|users
|array[link:#users[users]]
a|Specifies the list of users who belong to the group.
```

```
|===
```

.Example request

[%collapsible%closed]

====

```
[source,json,subs=+macros]
{
  "comment": "Admin group",
  "id": "5",
  "name": "Admin-Group",
  "policies": [
    {
      "name": "Policy1"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "name": "user-1"
    }
  ]
}
=====
```


== Response

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records

|records
|array[link:#s3_group[s3_group]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "comment": "Admin group",
      "id": "5",
      "name": "Admin-Group",
      "policies": [
        {
          "name": "Policy1"
        }
      ],
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "users": [
        {
          "name": "user-1"
        }
      ]
    }
  ]
}
```

```

    }
  ]
}
====

== Error

```

Status: Default

```

ONTAP Error Response Codes
//start table
[cols=2*,options=header]
|===
//header
| Error Code | Description
//end header
//end row
//start row
| 92405934 +
//end row
//start row
| "Users \{users} specified in the user list do not exist for SVM
"{svm.name}". Use the "object-store-server user create" command to create
a user.";
//end row
//start row
| 92405937 +
//end row
//start row
| "Policies \{policies} specified in the policy list do not exist for SVM
"{svm.name}". Use the "object-store-server policy create" command to
create a policy.";
//end row
|===
//end table

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]

```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#policies]
```

```
[.api-collapsible-fifth-title]
```

```
policies
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#users]

[.api-collapsible-fifth-title]

users

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|===

[#s3_group]

[.api-collapsible-fifth-title]

s3_group

This is a container for S3 user groups.

[cols=3*,options=header]

|===

|Name

|Type

|Description

```

|comment
|string
a|Can contain any additional information about the group being created or
modified.

|id
|integer
a|Specifies a unique group ID used to identify a particular group. This
parameter should not be specified in the POST method. A group ID is
automatically generated and it is retrieved using the GET method. Group id
is SVM scoped.

|name
|string
a|Specifies the name of the group. A group name length can range from 1 to
128 characters and can only contain the following combination of
characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|policies
|array[link:#policies[policies]]
a|Specifies a list of policies that are attached to the group. The
wildcard character "*" is a valid value for specifying all policies.

|svm
|link:#svm[svm]
a|

|users
|array[link:#users[users]]
a|Specifies the list of users who belong to the group.

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

|===

//end collapsible .Definitions block

====

[[IDfe116dcb616b9e19f75cf694c1a94ec3]]

= Delete the S3 group configuration for an SVM

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/protocols/s3/services/{svm.uuid}/groups/{id}`#

Introduced In: 9.8

Deletes the S3 group configuration of an SVM.

== Related ONTAP commands

* `vserver object-store-server group delete`

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_groups_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/groups]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|id

|string

|path

|True

a|Group identifier that identifies the unique group.

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
```



```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]

```

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[IDb514d22a38b777576059ef01cbf10dc6]]

= Retrieve the S3 group configuration for an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/s3/services/{svm.uuid}/groups/{id}`#

Introduced In: 9.8

Retrieves the S3 group configuration of an SVM.

== Related ONTAP commands

* `vserver object-store-server group show`

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_groups_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/groups]

```

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|id
|string
|path
|True
a|Group identifier that identifies the unique group.

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Can contain any additional information about the group being created or

```

modified.

|id

|integer

a|Specifies a unique group ID used to identify a particular group. This parameter should not be specified in the POST method. A group ID is automatically generated and it is retrieved using the GET method. Group id is SVM scoped.

|name

|string

a|Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|policies

|array[link:#policies[policies]]

a|Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies.

|svm

|link:#svm[svm]

a|

|users

|array[link:#users[users]]

a|Specifies the list of users who belong to the group.

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "comment": "Admin group",
  "id": "5",
  "name": "Admin-Group",
  "policies": [
    {
      "_links": {
```

```

        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "Policy1"
}
],
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"users": [
    {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "user-1"
    }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#policies]
[.api-collapsible-fifth-title]
policies

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Specifies the name of the policy. A policy name length can range from 1
to 128 characters and can only contain the following combination of
characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#users]
[.api-collapsible-fifth-title]
users
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```



```
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```

|===

//end collapsible .Definitions block
=====

[[IDaa0693190c95c53a00da307c615a9cbd]]
= Update S3 group configuration for an SVM

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/groups/{id}`#

*Introduced In:* 9.8

Updates the S3 group configuration of an SVM.

== Important notes

* The following fields can be modified for a group:
* `name` - Group name that needs to be modified.
* `users` - List of users present in the group.
* `policies` - List of policies to be attached to this group.

== Recommended optional properties

* `comment` - Short description about the S3 Group.

== Related ONTAP commands

* `vserver object-store-server group modify`

== Learn more

*
xref:{relative_path}protocols_s3_services_svm.uuid_groups_endpoint_overvie
w.html[DOC /protocols/s3/services/{svm.uuid}/groups]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type

```

```
|In
|Required
|Description

|id
|integer
|path
|True
a|Group identifier that identifies the unique group.
```

```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|comment
|string
a|Can contain any additional information about the group being created or modified.
```

```
|id
|integer
a|Specifies a unique group ID used to identify a particular group. This parameter should not be specified in the POST method. A group ID is automatically generated and it is retrieved using the GET method. Group id is SVM scoped.
```

```
|name
|string
a|Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of
```

characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

```
|policies
|array[link:#policies[policies]]
a|Specifies a list of policies that are attached to the group. The
wildcard character "*" is a valid value for specifying all policies.
```

```
|svm
|link:#svm[svm]
a|
```

```
|users
|array[link:#users[users]]
a|Specifies the list of users who belong to the group.
```

```
|===
```

```
.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "comment": "Admin group",
  "id": "5",
  "name": "Admin-Group",
  "policies": [
    {
      "name": "Policy1"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "users": [
    {
      "name": "user-1"
    }
  ]
}
=====
```

```
== Response
```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

//start table

[cols=2*,options=header]

|==

//header

| Error Code | Description

//end header

//end row

//start row

|92405934 +

//end row

//start row

|"Users \{users} specified in the user list do not exist for SVM
"{svm.name}". Use the "object-store-server user create" command to create
a user.";

//end row

//start row

|92405937 +

//end row

//start row

|"Policies \{policies} specified in the policy list do not exist for SVM
"{svm.name}". Use the "object-store-server policy create" command to
create a policy.";

//end row

|==

//end table

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#href]

[.api-collapsible-fifth-title]

href

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#policies]
```

```
[.api-collapsible-fifth-title]
```

```
policies
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#users]
[.api-collapsible-fifth-title]
users
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|Specifies the name of the user. A user name length can range from 1 to
64 characters and can only contain the following combination of characters
0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".
```

```
|===
```

```
[#s3_group]
[.api-collapsible-fifth-title]
s3_group
```

This is a container for S3 user groups.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|comment
|string
```

a|Can contain any additional information about the group being created or modified.

|id

|integer

a|Specifies a unique group ID used to identify a particular group. This parameter should not be specified in the POST method. A group ID is automatically generated and it is retrieved using the GET method. Group id is SVM scoped.

|name

|string

a|Specifies the name of the group. A group name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|policies

|array[link:#policies[policies]]

a|Specifies a list of policies that are attached to the group. The wildcard character "*" is a valid value for specifying all policies.

|svm

|link:#svm[svm]

a|

|users

|array[link:#users[users]]

a|Specifies the list of users who belong to the group.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description


```

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

```

====

[[ID89cb7e79fdde492c2cba905dd647da95]]

= Retrieve S3 protocol historical performance metrics

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/s3/services/{svm.uuid}/metrics`#

Introduced In: 9.8

Retrieves historical performance metrics for the S3 protocol of an SVM.

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|timestamp

|string

|query

|False

a|Filter by timestamp

|duration

|string

|query

|False

a|Filter by duration

|throughput.total

|integer

|query

|False

a|Filter by throughput.total

```
|throughput.read
|integer
|query
|False
a|Filter by throughput.read

|throughput.other
|integer
|query
|False
a|Filter by throughput.other

|throughput.write
|integer
|query
|False
a|Filter by throughput.write

|latency.total
|integer
|query
|False
a|Filter by latency.total

|latency.read
|integer
|query
|False
a|Filter by latency.read

|latency.other
|integer
|query
|False
a|Filter by latency.other

|latency.write
|integer
|query
|False
a|Filter by latency.write
```

```
|status
|string
|query
|False
a|Filter by status
```

```
|iops.total
|integer
|query
|False
a|Filter by iops.total
```

```
|iops.read
|integer
|query
|False
a|Filter by iops.read
```

```
|iops.other
|integer
|query
|False
a|Filter by iops.other
```

```
|iops.write
|integer
|query
|False
a|Filter by iops.write
```

```
|svm.uuid
|string
|path
|True
a|Unique identifier of the SVM.
```

```
|interval
|string
|query
```

```

|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:

* 1h: Metrics over the most recent hour sampled over 15 seconds.
* 1d: Metrics over the most recent day sampled over 5 minutes.
* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#records[records]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
```

```

    "href": "/api/resourcelink"
  },
  "self": {
    "href": "/api/resourcelink"
  }
},
"records": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string

```



```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#iops]
[.api-collapsible-fifth-title]
iops

```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#records]
[.api-collapsible-fifth-title]
records
```

Performance numbers, such as IOPS latency and throughput.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|duration
|string
```

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

```
|iops
|link:#iops[iops]
```

a|The rate of I/O operations observed at the storage object.

```
|latency
|link:#latency[latency]
```

a|The round trip latency in microseconds observed at the storage object.

```

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage S3 service policies

```

```
:leveloffset: +1
```

```
[[IDf11765edfa8ccbc1b5fc6f2e1c813d54]]  
= Protocols S3 services svm.uuid policies endpoint overview
```

```
== Overview
```

An S3 policy is an object that when associated with a resource, defines their permissions. Buckets and objects are defined as resources. Policies are used to manage access to these resources.

```
== Examples
```

```
=== Retrieving all fields for all S3 policies of an SVM
```

```
----
```

```
# The API:
```

```
/api/protocols/s3/services/{svm.uuid}/policies:
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-  
11e9-8c06-  
0050568ea123/policies?fields=*&return_records=true&return_timeout=15" -H  
"accept: application/json"
```

```
# The response:
```

```
{  
  "num_records": 2,  
  "records": [  
    {  
      "comment": "S3 policy.",  
      "name": "Policy1",  
      "statements": [  
        {  
          "actions": [  
            "*"   
          ],  
          "effect": "allow",  
          "index": 0,  
          "resources": [  
            "bucket1",  
            "bucket1/*"  
          ],  
        }  
      ]  
    }  
  ]  
}
```

```

        "sid": "FullAccessToBucket1"
    },
    {
        "actions": [
            "DeleteObject"
        ],
        "effect": "deny",
        "index": 1,
        "resources": [
            "*"
        ],
        "sid": "DenyDeleteObjectAccessToAllResources"
    }
],
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
},
{
    "comment": "S3 policy 2.",
    "name": "Policy2",
    "statements": [
        {
            "actions": [
                "GetObject"
            ],
            "effect": "allow",
            "index": 3,
            "resources": [
                "*"
            ],
            "sid": "AllowGetObjectAccessToAllResources"
        },
        {
            "actions": [
                "*"
            ],
            "effect": "deny",
            "index": 3,
            "resources": [
                "*"
            ],
            "sid": "DenyAccessToAllResources"
        }
    ]
},

```



```

    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
  ]
}
-----

```

=== Retrieving the specified policy in the SVM

The API:

/api/protocols/s3/services/{svm.uuid}/policies/{name}:

The call:

```

curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568ea123/policies/Policy1?fields=*" -H "accept: application/json"

```

The response:

```

{
  "comment": "S3 policy.",
  "name": "Policy1",
  "statements": [
    {
      "actions": [
        "GetObject",
        "PutObject",
        "DeleteObject",
        "ListBucket",
        "ListMyBuckets"
      ],
      "effect": "deny",
      "index": 0,
      "resources": [
        "*"
      ],
      "sid": "DenyAccessToAllResources"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}

```

```

}
----

=== Creating an S3 policy for an SVM

----

# The API:
/api/protocols/s3/services/{svm.uuid}/policies

# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-11e9-8c06-0050568eal23/policies?return_records=true" -H "accept: application/json" -H "authorization: Basic YWRtaW46bmV0YXBwMSE=" -H "Content-Type: application/json" -d "{ \"comment\": \"S3 policy.\", \"name\": \"Policy1\", \"statements\": [ { \"actions\": [ \"ListBucket\", \"ListMyBuckets\" ], \"effect\": \"allow\", \"resources\": [ \"*\" ], \"sid\": \"AllowListAccessToAllResources\" } ]}"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 289
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "comment": "S3 policy.",
      "name": "Policy1",
      "statements": [
        {
          "actions": [
            "ListBucket",
            "ListMyBuckets"
          ],
          "effect": "allow",
          "index": 5,
          "resources": [
            "*"
          ],
          "sid": "AllowListAccessToAllResources"
        }
      ]
    }
  ]
}

```

```

    ],
    "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
}
]
}
-----

```

=== Updating an S3 policy for an SVM

The API:

/api/protocols/s3/services/{svm.uuid}/policies/{name}:

The call:

```

curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-
11e9-8c06-0050568ea123/policies/Policy1" -H "accept: application/json" -H
"authorization: Basic YWRtaW46bmV0YXBwMSE=" -H "Content-Type:
application/json" -d "{ \"comment\": \"S3 policy.\", \"statements\": [ {
\"actions\": [ \"GetObject\", \"PutObject\", \"DeleteObject\",
\"ListBucket\", \"ListMyBuckets\" ], \"effect\": \"allow\", \"resources\":
[ \"bucket1\", \"bucket1/*\" ], \"sid\": \"FullAccessToAllResources\" }
] }"

```

The response:

```

HTTP/1.1 202 Accepted
Date: Mon, 16 Mar 2020 11:32:27 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
-----

```

=== Deleting an S3 policy for a specified SVM

The API:

/api/protocols/s3/services/{svm.uuid}/policies/{name}:

The call:

```

curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/12f3ba4c-7ae0-

```

```
11e9-8c06-0050568ea123/policies/Policy1" -H "accept: application/json"
```

```
# The response:
```

```
HTTP/1.1 202 Accepted
```

```
Date: Mon, 16 Mar 2020 11:32:27 GMT
```

```
Server: libzapid-httpd
```

```
X-Content-Type-Options: nosniff
```

```
Cache-Control: no-cache,no-store,must-revalidate
```

```
Content-Length: 189
```

```
Content-Type: application/json
```

```
----
```

```
[[ID8ec2d105a8b99a72a5cdd9cf98392d8e]]
```

```
= Retrieve S3 policies in SVM configuration
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/protocols/s3/services/{svm.uuid}/policies`#
```

```
*Introduced In:* 9.8
```

```
Retrieves the S3 policies SVM configuration.
```

```
== Related ONTAP commands
```

```
* `vserver object-store-server policy show`
```

```
== Learn more
```

```
*
```

```
xref:{relative_path}protocols_s3_services_svm.uuid_policies_endpoint_overv  
iew.html[DOC /protocols/s3/services/{svm.uuid}/policies]
```

```
== Parameters
```

```
[cols=5*,options=header]  
|==
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|name  
|string  
|query  
|False  
a|Filter by name
```

```
|statements.sid  
|string  
|query  
|False  
a|Filter by statements.sid
```

```
|statements.effect  
|string  
|query  
|False  
a|Filter by statements.effect
```

```
|statements.index  
|integer  
|query  
|False  
a|Filter by statements.index
```

```
|statements.resources  
|string  
|query  
|False  
a|Filter by statements.resources
```

```
|statements.actions  
|string  
|query  
|False  
a|Filter by statements.actions
```

```
|comment  
|string  
|query  
|False
```

a|Filter by comment

|svm.uuid

|string

|query

|False

a|Filter by svm.uuid

|svm.name

|string

|query

|False

a|Filter by svm.name

|read-only

|boolean

|query

|False

a|Filter by read-only

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|return_records

|boolean

```

|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

```

== Response

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records

```

```

|integer
a|Number of records

|records
|array[link:#s3_policy[s3_policy]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "comment": "S3 policy.",
      "name": "Policy1",
      "statements": [
        {
          "actions": [
            "*"
          ],
          "effect": "allow",
          "index": 0,
          "resources": [
            "bucket1",
            "bucket1/*"
          ],
          "sid": "FullAccessToBucket1"
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    }
  ]
}

```



```

    }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

```

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#s3_policy_statement]
[.api-collapsible-fifth-title]

```

s3_policy_statement

Specifies information about a single access policy statement.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|actions

|array[string]

a|For each resource, S3 supports a set of operations. The resource operations allowed or denied are identified by an action list:

* GetObject - retrieves objects from a bucket.

* PutObject - puts objects in a bucket.

* DeleteObject - deletes objects from a bucket.

* ListBucket - lists the objects in a bucket.

* GetBucketAcl - retrieves the access control list (ACL) of a bucket.

* GetObjectAcl - retrieves the access control list (ACL) of an object.

* ListAllMyBuckets - lists all of the buckets in a server.

The wildcard character "*" can be used to form a regular expression for specifying actions.

|effect

|string

a|Specifies whether access is allowed or denied. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.

|index

|integer

a|Specifies a unique statement index used to identify a particular statement. This parameter should not be specified in the POST method. A statement index is automatically generated and is retrieved using the GET method.

|resources

|array[string]

a|

```

|sid
|string
a|Specifies the statement identifier which contains additional information
about the statement.

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string

```

a|The unique identifier of the SVM.

|===

[#s3_policy]
[.api-collapsible-fifth-title]
s3_policy

An S3 policy is an object. It defines resource (bucket, folder or object) permissions. These policies get evaluated when an object store user user makes a request. Permissions in the policies determine whether the request is allowed or denied.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string

a|Can contain any additional information about the S3 policy.

|name
|string

a|Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|read-only
|boolean

a|Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method.

|statements
|array[link:#s3_policy_statement[s3_policy_statement]]
a|Specifies the policy statements.

|svm
|link:#svm[svm]

```

a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID7889e42eda3458193523aed4b550d4d5]]
= Create an S3 policy configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/policies`#

*Introduced In:* 9.8

Creates the S3 policy configuration.

== Important notes

* Each SVM can have one or more s3 policy configurations.

== Required properties

* `svm.uuid` - Existing SVM in which to create the s3 policy
configuration.
* `name` - Policy name that is to be created.

== Recommended optional properties

* `comment` - Short description about the S3 policy.
* `statements.effect` - Indicates whether to allow or deny access.
* `statements.actions` - List of actions that can be allowed or denied
access. Example: GetObject, PutObject, DeleteObject, ListBucket and
ListMyBuckets.
* `statements.resources` - Buckets or objects that can be allowed or

```

denied access.

* `statements.sid` - Statement identifier providing additional information about the statement.

== Related ONTAP commands

* `vserver object-store-server policy create`

* `vserver object-store-server policy add-statement`

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_policies_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/policies]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|===

== Request Body


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Can contain any additional information about the S3 policy.

|name
|string
a|Specifies the name of the policy. A policy name length can range from 1
to 128 characters and can only contain the following combination of
characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|read-only
|boolean
a|Specifies whether or not the s3 policy is read only. This parameter
should not be specified in the POST method.

|statements
|array[link:#s3_policy_statement[s3_policy_statement]]
a|Specifies the policy statements.

|svm
|link:#svm[svm]
a|

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "comment": "S3 policy.",
  "name": "Policy1",
  "statements": [
    {
      "actions": [

```

```

        "*"
    ],
    "effect": "allow",
    "index": 0,
    "resources": [
        "bucket1",
        "bucket1/*"
    ],
    "sid": "FullAccessToBucket1"
}
],
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Response

```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records

|records
|array[link:#s3_policy[s3_policy]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
    "records": [

```

```

{
  "comment": "S3 policy.",
  "name": "Policy1",
  "statements": [
    {
      "actions": [
        "*"
      ],
      "effect": "allow",
      "index": 0,
      "resources": [
        "bucket1",
        "bucket1/*"
      ],
      "sid": "FullAccessToBucket1"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

Error Code	Description
-----	-----

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

====

[#s3_policy_statement]

[.api-collapsible-fifth-title]

s3_policy_statement

Specifies information about a single access policy statement.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|actions

|array[string]

a|For each resource, S3 supports a set of operations. The resource operations allowed or denied are identified by an action list:

* GetObject - retrieves objects from a bucket.

* PutObject - puts objects in a bucket.

* DeleteObject - deletes objects from a bucket.

* ListBucket - lists the objects in a bucket.

* GetBucketAcl - retrieves the access control list (ACL) of a bucket.

* GetObjectAcl - retrieves the access control list (ACL) of an object.

* ListAllMyBuckets - lists all of the buckets in a server.

The wildcard character "*" can be used to form a regular expression for specifying actions.

|effect

|string

a|Specifies whether access is allowed or denied. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.

|index

|integer

a|Specifies a unique statement index used to identify a particular statement. This parameter should not be specified in the POST method. A statement index is automatically generated and is retrieved using the GET method.

|resources

|array[string]

a|

```

|sid
|string
a|Specifies the statement identifier which contains additional information
about the statement.

|===

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

|===

```
[#s3_policy]
[.api-collapsible-fifth-title]
s3_policy
```

An S3 policy is an object. It defines resource (bucket, folder or object) permissions. These policies get evaluated when an object store user user makes a request. Permissions in the policies determine whether the request is allowed or denied.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|comment
```

```
|string
```

a|Can contain any additional information about the S3 policy.

```
|name
```

```
|string
```

a|Specifies the name of the policy. A policy name length can range from 1 to 128 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

```
|read-only
```

```
|boolean
```

a|Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method.

```
|statements
```

```
|array[link:#s3_policy_statement[s3_policy_statement]]
```

a|Specifies the policy statements.

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

```
[#collection_links]  
[.api-collapsible-fifth-title]  
collection_links  
[#error_arguments]  
[.api-collapsible-fifth-title]  
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|arguments  
|array[link:#error_arguments[error_arguments]]  
a|Message arguments
```

```
|code  
|string
```

```

a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID9e70bfc352a6950a203a7e6d67fa07fe]]
= Delete an S3 policy configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/policies/{name}`#

*Introduced In:* 9.8

Deletes the S3 policy configuration of an SVM.

== Related ONTAP commands

* `vserver object-store-server policy delete`
* `vserver object-store-server policy delete-statement`

== Learn more

*
xref:{relative_path}protocols_s3_services_svm.uuid_policies_endpoint_overv
iew.html[DOC /protocols/s3/services/{svm.uuid}/policies]

== Parameters

[cols=5*,options=header]
|===

```



```
|Name
|Type
|In
|Required
|Description

|name
|string
|path
|True
a|Policy name

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
```

```

====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]

```

```

|===

```

```

|Name

```

```

|Type

```

```

|Description

```

```

|code

```

```

|string

```

```

a|Argument code

```

```

|message

```

```

|string

```

```

a|Message argument

```

```

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID99eb020faed68adaebea5eca0bac2d5c]]
= Retrieve an S3 policy configuration

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/policies/{name}`#

*Introduced In:* 9.8

Retrieves the S3 policy configuration of an SVM.

```

== Related ONTAP commands

* `vserver object-store-server policy show`

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_policies_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/policies]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|name

|string

|path

|True

a|Policy name

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|comment
|string
a|Can contain any additional information about the S3 policy.

|name
|string
a|Specifies the name of the policy. A policy name length can range from 1
to 128 characters and can only contain the following combination of
characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|read-only
|boolean
a|Specifies whether or not the s3 policy is read only. This parameter
should not be specified in the POST method.

|statements
|array[link:#s3_policy_statement[s3_policy_statement]]
a|Specifies the policy statements.

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "comment": "S3 policy.",
  "name": "Policy1",
  "statements": [
```

```

{
  "actions": [
    "*"
  ],
  "effect": "allow",
  "index": 0,
  "resources": [
    "bucket1",
    "bucket1/*"
  ],
  "sid": "FullAccessToBucket1"
}
],
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====

```

```
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#s3_policy_statement]
[.api-collapsible-fifth-title]
s3_policy_statement

Specifies information about a single access policy statement.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|actions
|array[string]
a|For each resource, S3 supports a set of operations. The resource
operations allowed or denied are identified by an action list:

* GetObject - retrieves objects from a bucket.
* PutObject - puts objects in a bucket.
* DeleteObject - deletes objects from a bucket.
* ListBucket - lists the objects in a bucket.
* GetBucketAcl - retrieves the access control list (ACL) of a bucket.
```

* GetObjectAcl - retrieves the access control list (ACL) of an object.
* ListAllMyBuckets - lists all of the buckets in a server.
The wildcard character "*" can be used to form a regular expression for specifying actions.

|effect
|string
a|Specifies whether access is allowed or denied. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.

|index
|integer
a|Specifies a unique statement index used to identify a particular statement. This parameter should not be specified in the POST method. A statement index is automatically generated and is retrieved using the GET method.

|resources
|array[string]
a|

|sid
|string
a|Specifies the statement identifier which contains additional information about the statement.

|===

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string


```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[IDad6895c52afdf6eac4c3cef52d5b3239]]

= Update an S3 policy configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-block]#`/protocols/s3/services/{svm.uuid}/policies/{name}`#

Introduced In: 9.8

Updates the S3 policy configuration of an SVM.

== Important notes

* The following fields can be modified for a policy:

*** `comment` - Any information related to the policy.

*** `statements` - Specifies the array of policy statements.

== Related ONTAP commands

* `vserver object-store-server policy modify`

* `vserver object-store-server policy modify-statement`

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_policies_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/policies]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|name
```

```
|string
```

```
|path
```

```
|True
```

```
a|Policy name
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|UUID of the SVM to which this object belongs.
```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|comment
```

```
|string
```

```
a|Can contain any additional information about the S3 policy.
```

```
|read-only
```

```
|boolean
```

```
a|Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method.
```

```
|statements
```

```
|array[link:#s3_policy_statement[s3_policy_statement]]
```

a|Specifies the policy statements.

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "comment": "S3 policy.",
  "statements": [
    {
      "actions": [
        "*"
      ],
      "effect": "allow",
      "index": 0,
      "resources": [
        "bucket1",
        "bucket1/*"
      ],
      "sid": "FullAccessToBucket1"
    }
  ],
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

====

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
------------	-------------

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#s3_policy_statement]
[.api-collapsible-fifth-title]
s3_policy_statement
```

Specifies information about a single access policy statement.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|actions
|array[string]
```

a|For each resource, S3 supports a set of operations. The resource operations allowed or denied are identified by an action list:

- * GetObject - retrieves objects from a bucket.
- * PutObject - puts objects in a bucket.
- * DeleteObject - deletes objects from a bucket.
- * ListBucket - lists the objects in a bucket.
- * GetBucketAcl - retrieves the access control list (ACL) of a bucket.
- * GetObjectAcl - retrieves the access control list (ACL) of an object.
- * ListAllMyBuckets - lists all of the buckets in a server.

The wildcard character "*" can be used to form a regular expression for specifying actions.

```
|effect
|string
```

a|Specifies whether access is allowed or denied. If access (to allow) is not granted explicitly to a resource, access is implicitly denied. Access

can also be denied explicitly to a resource, in order to make sure that a user cannot access it, even if a different policy grants access.

|index

|integer

a|Specifies a unique statement index used to identify a particular statement. This parameter should not be specified in the POST method. A statement index is automatically generated and is retrieved using the GET method.

|resources

|array[string]

a|

|sid

|string

a|Specifies the statement identifier which contains additional information about the statement.

|===

[#href]

[.api-collapsible-fifth-title]

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

|string

a|

|===

[#_links]

[.api-collapsible-fifth-title]

_links

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#s3_policy]

[.api-collapsible-fifth-title]

s3_policy

An S3 policy is an object. It defines resource (bucket, folder or object) permissions. These policies get evaluated when an object store user user makes a request. Permissions in the policies determine whether the request is allowed or denied.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|comment

|string

a|Can contain any additional information about the S3 policy.

|read-only

|boolean

a|Specifies whether or not the s3 policy is read only. This parameter should not be specified in the POST method.


```
|statements
|array[link:#s3_policy_statement[s3_policy_statement]]
a|Specifies the policy statements.
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage S3 service users
```

```
:leveloffset: +1
```

```
[[ID4e3e458fd9b3a8417341806a21eda662]]
```

```
= Protocols S3 services svm.uuid users endpoint overview
```

```
== Overview
```

An S3 user account is created on the S3 server. Buckets that are created for the server are associated with that user (as the owner of the buckets).

The creation of the user account involves generating a pair of keys "access" and "secret".

These keys are shared with clients (by the administrator out of band) who want to access the S3 server. The `access_key` is sent in the request and it identifies the user performing the operation. The client or server never send the `secret_key` over the wire.

Only the `access_key` can be retrieved from a GET operation. The `secret_key` along with the `access_key` is returned from a POST operation and from a PATCH operation if the administrator needs to regenerate the keys.

== Examples

=== Retrieving S3 user configurations for a particular SVM

The API:

/api/protocols/s3/services/{svm.uuid}/users

The call:

```
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-11e9-99e1-0050568e3ed9/users?fields=*&return_records=true" -H "accept: application/hal+json"
```

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "db2ec036-8375-11e9-99e1-0050568e3ed9",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/db2ec036-8375-11e9-99e1-0050568e3ed9"
          }
        }
      },
      "name": "user-1",
      "comment": "S3 user",
      "access_key":
      "8OP1Yd5gm53sTNkTNgrsJ0_4iHvw_Ir_9xtDhzGa3m2_a_Yhtv6Bm3Dq_Xv79Stq90Bwa5NrTL7UQ2u_0xN0IW_x39cm1h3sn69fN6cf6STA48W05PAxuGED3NcR7rsn",
      "_links": {
        "self": {
          "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-0050568e3ed9/users/user-1"
        }
      }
    }
  ]
}
```

```

    },
    {
      "svm": {
        "uuid": "db2ec036-8375-11e9-99e1-0050568e3ed9",
        "name": "vs1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/db2ec036-8375-11e9-99e1-0050568e3ed9"
          }
        }
      },
      "name": "user-2",
      "comment": "s3-user",
      "access_key":
        "uYo34d4eR8a3is7JDSCY1xrNwL7gFMA338ZEX2mNrgJ34Kb4u98QNhBGT3ghs9GA2bzNdYBSn
        5_rBfjIY4mt36CMFE4d3g0L3Pa_2nXD6g6CAq_D0422LK__pbH6wvy8",
      "_links": {
        "self": {
          "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
          0050568e3ed9/users/user-2"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
      0050568e3ed9/users?fields=*amp;return_records=true"
    }
  }
}
----

```

=== Retrieving the user configuration of a specific S3 user

The API:

/api/protocols/s3/services/{svm.uuid}/users/{name}

The call:

```
curl -X GET "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users/user-1" -H "accept: application/hal+json"
```

The response:

```
{
  "svm": {
    "uuid": "db2ec036-8375-11e9-99e1-0050568e3ed9",
    "name": "vs1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/db2ec036-8375-11e9-99e1-0050568e3ed9"
      }
    }
  },
  "name": "user-1",
  "comment": "s3-user",
  "access_key":
    "uYo34d4eR8a3is7JDSCY1xrNwL7gFMA338ZEX2mNrgJ34Kb4u98QNhBGT3ghs9GA2bzNdYBSn
    5_rBfjIY4mt36CMFE4d3g0L3Pa_2nXD6g6CAq_D0422LK__pbH6wvy8",
  "_links": {
    "self": {
      "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-1"
    }
  }
}
```

=== Creating an S3 user configuration

```
# The API:
/api/protocols/s3/services/{svm.uuid}/users
```

```
# The call:
curl -X POST "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users" -H "accept: application/json" -H "Content-
Type: application/json" -d '{"name": "user-1"}'
```

```
# The response:
HTTP/1.1 201 Created
Date: Fri, 31 May 2019 09:34:25 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-1
Content-Length: 244
Content-Type: application/json
```

```
{
  "num_records": 1,
  "records": [
    {
      "name": "user-1",
      "access_key":
"8OP1Yd5gm53sTNkTNgrsJ0_4iHvw_Ir_9xtDhzGa3m2_a_Yhtv6Bm3Dq_Xv79Stq90Bwa5NrT
L7UQ2u_0xN0IW_x39cm1h3sn69fN6cf6STA48W05PAxuGED3NcR7rsn",
      "secret_key":
"SSS4oNA7_43yfu_zs938T5nY9xYZccFq_60_Q925h4t535km313qb0bDvdQ2MIK_8ebVf0gnD
06K8qcNBg3t_KcpjHTXA2elshTEjrdMhsM9b47uOdQGw4Mex6yrbPgr"
    }
  ]
}
----
```

=== Regenerating keys for a specific S3 user for the specified SVM

The API:

/api/protocols/s3/services/{svm.uuid}/users/{name}

The call:

```
curl -X PATCH "https://<mgmt-ip>/api/protocols/s3/services/db2ec036-8375-
11e9-99e1-0050568e3ed9/users/user-2?regenerate_keys=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ }"
```

The response:

HTTP/1.1 200 OK

Date: Fri, 31 May 2019 09:55:45 GMT

Server: libzapid-httpd

X-Content-Type-Options: nosniff

Cache-Control: no-cache,no-store,must-revalidate

Content-Length: 391

Content-Type: application/hal+json

```
{
  "num_records": 1,
  "records": [
    {
      "name": "user-2",
      "access_key":
"hUod3l_sg632PjPlTgdQNKWXI3E_yTra0h96xrpsAPly3Qa_KmYYXq3kIuAJ3CyD4gVOakjj_
PwVIVjATP1C2t1IQ3KB_9ctS1Ph921b1C17N6Y0PtWfv6AZD__j_C4j",
      "secret_key":
"3w03fT_7Pv328_dYB8FN4YsD101Hn0ilu_gmqOenYydaNc22c7AIDN46c__T_5y0A3Y69w412
```

```
F13A1bzJSpXH4C0nNAP4N_Ce1_Z_9_d7bA08bs28ccw50ab_4osA3bq",
  "_links": {
    "self": {
      "href": "/api/protocols/s3/services/db2ec036-8375-11e9-99e1-
0050568e3ed9/users/user-2"
    }
  }
}
]
}
----
```

=== Deleting the specified S3 user configuration for a specified SVM

The API:

/api/protocols/s3/services/{svm.uuid}/users/{name}

The call:

```
curl -X DELETE "https://<mgmt-ip>/api/protocols/s3/services/03ce5c36-f269-
11e8-8852-0050568e5298/users/user-2" -H "accept: application/json"
```

```
[[IDaa431befcdf87d9ddd587204e4ec09c1]]
```

= Retrieve the SVM configuration for an S3 user

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/users`#
```

Introduced In: 9.7

Retrieves the S3 user's SVM configuration.

== Related ONTAP commands

* `vserver object-store-server user show`

== Learn more

*

```
xref:{relative_path}protocols_s3_services_svm.uuid_users_endpoint_overview
.html[DOC /protocols/s3/services/{svm.uuid}/users]
```

```

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|access_key
|string
|query
|False
a|Filter by access_key

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|comment
|string
|query
|False
a|Filter by comment

|name
|string
|query
|False
a|Filter by name

```



```
|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1
* Max value: 120
* Min value: 0

```
|order_by
```

```
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#collection_links[collection_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#s3_user[s3_user]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  }
}
```

```

},
"records": [
  {
    "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
    "comment": "S3 user",
    "name": "user-1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {

```

```

    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```

```

[#collection_links]
[.api-collapsible-fifth-title]
collection_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#s3_user]
[.api-collapsible-fifth-title]
s3_user
```

This is a container of S3 users.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|access_key
|string
a|Specifies the access key for the user.
```

```
|comment
|string
a|Can contain any additional information about the user being created or
modified.
```

```
|name
|string
a|Specifies the name of the user. A user name length can range from 1 to
64 characters and can only contain the following combination of characters
0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
[#error_arguments]
```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

```

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID497518c1f143b64b858124e2f61b138e]]
= Create an S3 user configuration

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/users`#

*Introduced In:* 9.7

Creates the S3 user configuration.

== Important notes

* Each SVM can have one or more user configurations.
* If user creation is successful, a user access_key and secret_key is
returned as part of the response.

== Required properties

* `svm.uuid` - Existing SVM in which to create the user configuration.
* `name` - User name that is to be created.

== Default property values

* `comment` - ""

== Related ONTAP commands

* `vserver object-store-server user create`

== Learn more

*
xref:{relative_path}protocols_s3_services_svm.uuid_users_endpoint_overview
.html[DOC /protocols/s3/services/{svm.uuid}/users]

```


== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|svm.uuid

|string

|path

|True

a|UUID of the SVM to which this object belongs.

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|access_key

|string

a|Specifies the access key for the user.

|comment

|string

a|Can contain any additional information about the user being created or modified.

|name

|string

a|Specifies the name of the user. A user name length can range from 1 to 64 characters and can only contain the following combination of characters 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|svm

|link:#svm[svm]

a|

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "access_key":

"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",

 "comment": "S3 user",

 "name": "user-1",

 "svm": {

 "name": "svm1",

 "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"

 }

}

====

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#s3_service_user_post_response[s3_service_user_post_response]]
a|

|===


.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
      "name": "user-1",
      "secret_key":
"A20_tDhC_cux2C2BmtL45bXB_a_Q65c_96FsAcOdo14Az8V31jBKDTc0uCL62Bh559gPB8s9r
rn0868QrF38_1dsV2u1_9H2tSf3qQ5xp9NT259C6z_GiZQ883Qn63X1"
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

```

```
| 92405773
| Object store server is not present for specified SVM. Create a object
store server and retry the operation.

| 92405787
| The specified user name contains invalid characters. Valid characters
for a user name are 0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

| 92405788
| User names must have between 1 and 64 characters.

| 92405791
| Failed to create access-key and secret-key.

| 92405817
| S3 users can be created only on data SVM.
|===
```

```
== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#s3_user]
[.api-collapsible-fifth-title]
s3_user
```

This is a container of S3 users.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|access_key
|string
a|Specifies the access key for the user.
```

```
|comment
|string
a|Can contain any additional information about the user being created or
modified.
```

```

|name
|string
a|Specifies the name of the user. A user name length can range from 1 to
64 characters and can only contain the following combination of characters
0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|svm
|link:#svm[svm]
a|

|===

[#collection_links]
[.api-collapsible-fifth-title]
collection_links
[#s3_service_user_post_response]
[.api-collapsible-fifth-title]
s3_service_user_post_response

[cols=3*,options=header]
|===
|Name
|Type
|Description

|access_key
|string
a|Specifies the access key for the user.

|name
|string
a|The name of the user.

|secret_key
|string
a|Specifies the secret key for the user.

|===

[#error_arguments]
[.api-collapsible-fifth-title]

```

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDabc72dcdaf2b9f7da3d4d20362665f2d]]
= Delete an S3 user configuration

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/users/{name}`#

*Introduced In:* 9.7

Deletes the S3 user configuration of an SVM.

== Related ONTAP commands

* `vserver object-store-server user delete`

== Learn more

*
xref:{relative_path}protocols_s3_services_svm.uuid_users_endpoint_overview
.html[DOC /protocols/s3/services/{svm.uuid}/users]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|name
|string
|path

```



```
|True
a|User name

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],

```

```

    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

```

```

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[ID0d876635f4030d3c4f7e0420db25f61e]]
= Retrieve the S3 user configuration for an SVM
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/users/{name}`#
```

Introduced In: 9.7

Retrieves the S3 user configuration of an SVM.

== Related ONTAP commands

* `vserver object-store-server user show`

== Learn more

*

xref:{relative_path}protocols_s3_services_svm.uuid_users_endpoint_overview.html[DOC /protocols/s3/services/{svm.uuid}/users]

```
== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|name
|string
|path
|True
a|User name

|svm.uuid
|string
|path
|True
a|UUID of the SVM to which this object belongs.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|access_key
```

```

|string
a|Specifies the access key for the user.

|comment
|string
a|Can contain any additional information about the user being created or
modified.

|name
|string
a|Specifies the name of the user. A user name length can range from 1 to
64 characters and can only contain the following combination of characters
0-9, A-Z, a-z, "_", "+", "=", ",", ".", "@", and "-".

|svm
|link:#svm[svm]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
  "comment": "S3 user",
  "name": "user-1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```



```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDeb30b333c20b95e7ebdb12f0ac19bbc9]]
= Update an S3 user configuration

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/s3/services/{svm.uuid}/users/{name}`#

*Introduced In:* 9.7

Updates the S3 user configuration of an SVM.

== Important notes

* User access_key and secret_key pair can be regenerated using the PATCH
operation.
* User access_key and secret_key is returned in a PATCH operation if the
"regenerate_keys" field is specified as true.

== Recommended optional properties

* `regenerate_keys` - Specifies if secret_key and access_key need to be
regenerated.
* `comment` - Any information related to the S3 user.

== Related ONTAP commands

```

```
* `vserver object-store-server user show`  
* `vserver object-store-server user regenerate-keys`
```

== Learn more

```
*  
xref:{relative_path}protocols_s3_services_svm.uuid_users_endpoint_overview  
.html[DOC /protocols/s3/services/{svm.uuid}/users]
```

== Parameters

```
[cols=5*,options=header]  
|==
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|name  
|string  
|path  
|True  
a|User name
```

```
|regenerate_keys  
|boolean  
|query  
|False  
a|Specifies whether or not to regenerate the user keys.
```

* Default value:

```
|svm.uuid  
|string  
|path  
|True  
a|UUID of the SVM to which this object belongs.
```

```
|==
```

== Request Body

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|access_key
|string
a|Specifies the access key for the user.

|comment
|string
a|Can contain any additional information about the user being created or
modified.

|svm
|link:#svm[svm]
a|

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
  "comment": "S3 user",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records


|records
|array[link:#s3_service_user_post_response[s3_service_user_post_response]]
a|

|===


.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "access_key":
"Pz3SB54G2B_6dsXQPrA5HrTPcf478qoAW6_Xx6qyqZ948AgZ_7YfCf_9nO87YoZmskxx3cq41
U2JAH2M3_fs321B4rkzS3a_oC5_8u7D8j_45N8OsBCBPWGD_1d_ccfq",
      "name": "user-1",
      "secret_key":
"A20_tDhC_cux2C2BmtL45bXB_a_Q65c_96FsAcOdo14Az8V31jBKDTc0uCL62Bh559gPB8s9r
rn0868QrF38_1dsV2u1_9H2tSf3qQ5xp9NT259C6z_GiZQ883Qn63X1"
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

```

```
| 92405792
| Failed to regenerate access-key and secret-key for user.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#s3_user]
[.api-collapsible-fifth-title]
s3_user
```

This is a container of S3 users.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|access_key
```

```
|string
```

```
a|Specifies the access key for the user.
```

```
|comment
```

```
|string
```

```
a|Can contain any additional information about the user being created or
modified.
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

```
[#collection_links]
[.api-collapsible-fifth-title]
collection_links
[#s3_service_user_post_response]
[.api-collapsible-fifth-title]
s3_service_user_post_response
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|access_key
|string
a|Specifies the access key for the user.
```

```
|name
|string
a|The name of the user.
```

```
|secret_key
|string
a|Specifies the secret key for the user.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

:leveloffset: -1

:leveloffset: -1

= SAN

:leveloffset: +1

```



```
[[ID872ec17405eebbacb306bf03c4539649]]  
= SAN overview
```

`== Overview`

The storage area network (SAN) endpoints and objects enable you to configure, provision, and manage SAN-related objects.

```
'''
```

`== Fibre Channel`

`=== Logins`

Fibre Channel logins represent connections, formed by Fibre Channel initiators, that have successfully logged in to ONTAP. This represents the Fibre Channel login on which higher-level protocols, such as Fibre Channel Protocol (FCP) and Non-Volatile Memory Express over Fibre Channel (NVMe over FC), rely.

The Fibre Channel logins REST API provides information about active Fibre Channel logins.

`=== WWPN Aliases`

A WWPN (world wide port name) is a unique 64-bit identifier for a Fibre Channel initiator. It is displayed as a 16-character hexadecimal value. SAN administrators may find it easier to identify Fibre Channel initiators using an alias, especially in larger SANs.

The WWPN alias REST API allows you to create, delete and discover aliases for WWPNs.

`=== Services`

A Fibre Channel Protocol (FCP) service defines the properties of the Fibre Channel Protocol target for an SVM. There can be at most one FCP service for a given SVM. An SVM's FCP service must be created before FCP initiators can login to the SVM.

The Fibre Channel Protocol (FCP) service REST API allows you to create, update, delete, and discover Fibre Channel Services for SVMs. Fibre Channel interfaces are the logical endpoints for Fibre Channel

network connections to an SVM.

'''

== iSCSI

=== Credentials

An iSCSI credentials object defines the authentication credentials to be used between an iSCSI initiator and ONTAP. It identifies an authentication type, user names, and the passwords that must be used to authenticate a specific initiator.

The iSCSI credentials REST API allows you to create, update, delete, and discover iSCSI credential objects.

=== Services

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can login to the SVM.

The iSCSI service REST API allows you to create, update, delete, and discover iSCSI services for SVMs.

=== Sessions

An iSCSI session consists of one or more TCP connections that link an iSCSI initiator with an iSCSI target. TCP connections can be added and removed from an iSCSI session by the iSCSI initiator. Across all TCP connections within an iSCSI session, an initiator sees one and the same target. After the connection is established, iSCSI control, data, and status messages are communicated over the session.

The iSCSI sessions REST API provides information about iSCSI initiators that have successfully logged in to ONTAP.

=== Learn More

* `_IP Interfaces_` found in the `_networking_` section. IP interfaces are the logical endpoints for iSCSI network connections to an SVM.

'''

== Initiator Groups

An initiator group (igroup) is a collection of Fibre Channel WWPNS (world

wide port names), iSCSI IQNs (qualified names), iSCSI EUIs (extended unique identifiers), or any combination of these, that identify host initiators.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, a network administrator creates an initiator group containing the hosts' initiator names, and then creates a LUN map that associates the initiator group with the LUN.

The initiator group REST API allows you to create, update, delete, and discover initiator groups. It also enables you to add and remove initiators that can access the target and associated LUNs.

== LUN Maps

A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the group's initiators are granted access to the LUN. The relationship between an initiator group and a LUN is many initiator groups to many LUNs.

The LUN map REST API allows you to create, delete, and discover LUN maps.

== LUNs

A LUN is the logical representation of storage in a storage area network (SAN).

The LUN REST API allows you to create, update, delete, and discover LUNs.

= Retrieve FC port information

:leveloffset: +1

[[ID3e56848fcb887a7e31015876d0ec4989]]
= Network FC logins endpoint overview

== Overview

Fibre Channel (FC) logins represent connections formed by FC initiators that have successfully logged in to ONTAP. This represents the FC login on

which higher-level protocols such as Fibre Channel Protocol and NVMe over FC (NVMe/FC) rely.

The Fibre Channel logins REST API provides information about active FC logins.

== Examples

=== Retrieving all FC logins

The API:

GET /api/network/fc/logins

The call:

curl -X GET "https://<mgmt-ip>/api/network/fc/logins" -H "accept: application/hal+json"

The response:

```
{
  "records": [
    {
      "svm": {
        "uuid": "056403da-83a7-4b13-bc78-6a93e8ea3596",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/056403da-83a7-4b13-bc78-6a93e8ea3596"
          }
        }
      },
      "interface": {
        "uuid": "01056403-1383-bc4b-786a-93e8ea35969d",
        "name": "lif1",
        "_links": {
          "self": {
            "href": "/api/network/fc/interfaces/01056403-1383-bc4b-786a-93e8ea35969d"
          }
        }
      },
      "initiator": {
        "wwpn": "8b:21:2f:07:00:00:00:00"
      },
      "_links": {
```

```

    "self": {
      "href": "/api/network/fc/logins/01056403-1383-bc4b-786a-93e8ea35969d/8b%3A21%3A2f%3A07%3A00%3A00%3A00%3A00"
    }
  },
  {
    "svm": {
      "uuid": "056403da-83a7-4b13-bc78-6a93e8ea3596",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/056403da-83a7-4b13-bc78-6a93e8ea3596"
        }
      }
    },
    "interface": {
      "uuid": "02056403-1383-bc4b-786a-93e8ea35969d",
      "name": "lif2",
      "_links": {
        "self": {
          "href": "/api/network/fc/interfaces/02056403-1383-bc4b-786a-93e8ea35969d"
        }
      }
    },
    "initiator": {
      "wwpn": "8c:21:2f:07:00:00:00:00"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/logins/02056403-1383-bc4b-786a-93e8ea35969d/8c%3A21%3A2f%3A07%3A00%3A00%3A00%3A00"
      }
    }
  },
  {
    "svm": {
      "uuid": "156403da-83a7-4b13-bc78-6a93e8ea3596",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/156403da-83a7-4b13-bc78-6a93e8ea3596"
        }
      }
    }
  },

```

```

    "interface": {
      "uuid": "03056403-1383-bc4b-786a-93e8ea35969d",
      "name": "lif3",
      "_links": {
        "self": {
          "href": "/api/network/fc/interfaces/00056403-1383-bc4b-786a-93e8ea35969d"
        }
      }
    },
    "initiator": {
      "wwpn": "8a:21:2f:07:00:00:00:00"
    },
    "_links": {
      "self": {
        "href": "/api/network/fc/logins/00056403-1383-bc4b-786a-93e8ea35969d/8a%3A21%3A2f%3A07%3A00%3A00%3A00%3A00"
      }
    }
  ],
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/network/fc/logins"
    }
  }
}
----

'''

```

=== Retrieving all FC logins with data protocol _fcp_ in SVM _svm1_

The `svm.name` and `protocol` query parameters are used to perform the query.

```

# The API:
GET /api/network/fc/logins

```

```

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/logins?svm.name=svm1&protocol=fcp" -H "accept: application/hal+json"

```

```
# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "056403da-83a7-4b13-bc78-6a93e8ea3596",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/056403da-83a7-4b13-bc78-6a93e8ea3596"
          }
        }
      },
      "interface": {
        "uuid": "01056403-1383-bc4b-786a-93e8ea35969d",
        "name": "lif2",
        "_links": {
          "self": {
            "href": "/api/network/fc/interfaces/01056403-1383-bc4b-786a-93e8ea35969d"
          }
        }
      },
      "initiator": {
        "wwpn": "8b:21:2f:07:00:00:00:00"
      },
      "protocol": "fcp",
      "_links": {
        "self": {
          "href": "/api/network/fc/logins/01056403-1383-bc4b-786a-93e8ea35969d/8b%3A21%3A2f%3A07%3A00%3A00%3A00%3A00"
        }
      }
    },
    {
      "svm": {
        "uuid": "056403da-83a7-4b13-bc78-6a93e8ea3596",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/056403da-83a7-4b13-bc78-6a93e8ea3596"
          }
        }
      },
      "interface": {
        "uuid": "02056403-1383-bc4b-786a-93e8ea35969d",

```

```

        "name": "lif3",
        "_links": {
            "self": {
                "href": "/api/network/fc/interfaces/02056403-1383-bc4b-786a-
93e8ea35969d"
            }
        },
        "initiator": {
            "wwpn": "8c:21:2f:07:00:00:00:00"
        },
        "protocol": "fcp",
        "_links": {
            "self": {
                "href": "/api/network/fc/logins/02056403-1383-bc4b-786a-
93e8ea35969d/8c%3A21%3A2f%3A07%3A00%3A00%3A00%3A00"
            }
        }
    ],
    "num_records": 2,
    "_links": {
        "self": {
            "href": "/api/network/fc/logins?svm.name=svml&protocol=fcp"
        }
    }
}
----

'''

```

=== Retrieving all FC logins for initiators belonging to `_igroup1_` and returning all of their properties

The ``igroups.name`` query parameter is used to perform the query. The ``fields`` query parameter is used to return all of the properties.

The API:

GET /api/network/fc/logins

The call:

```

curl -X GET "https://<mgmt-
ip>/api/network/fc/logins?igroups.name=igroup1&fields=*" -H "accept:
application/hal+json"

```



```
# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "056403da-83a7-4b13-bc78-6a93e8ea3596",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/056403da-83a7-4b13-bc78-6a93e8ea3596"
          }
        }
      },
      "interface": {
        "uuid": "01056403-1383-bc4b-786a-93e8ea35969d",
        "name": "lif2",
        "wwpn": "8b:21:2f:07:00:00:00:00",
        "_links": {
          "self": {
            "href": "/api/network/fc/interfaces/01056403-1383-bc4b-786a-93e8ea35969d"
          }
        }
      },
      "initiator": {
        "wwpn": "8b:21:2f:07:00:00:00:00",
        "wwnn": "95:21:2f:07:00:00:00:00"
      },
      "igroups": [
        {
          "uuid": "243bbb8a-46e9-4b2d-a508-a62dc93df9d1",
          "name": "igroup1",
          "_links": {
            "self": {
              "href": "/api/protocols/san/igroups/243bbb8a-46e9-4b2d-a508-a62dc93df9d1"
            }
          }
        }
      ],
      "port_address": "8aa53",
      "protocol": "fcp",
      "_links": {
        "self": {
          "href": "/api/network/fc/logins/01056403-1383-bc4b-786a-93e8ea35969d/8b%3A21%3A2f%3A07%3A00%3A00%3A00%3A00"
        }
      }
    }
  ]
}
```

```

    }
  }
},
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/fc/logins?igroups.name=igroup1&fields=*"
  }
}
}
}
-----

```

[[ID105a673b3f6af3ef763e509053a4472f]]

= Retrieve FC logins

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/network/fc/logins`#

Introduced In: 9.6

Retrieves FC logins.

== Related ONTAP commands

* `vserver fcp initiator show`

== Learn more

* SAN: xref:{relative_path}network_fc_logins_endpoint_overview.html[DOC /network/fc/logins]

* NVMe: xref:{relative_path}network_fc_logins_endpoint_overview.html[DOC /network/fc/logins]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

```

|Required
|Description

|protocol
|string
|query
|False
a|Filter by protocol

|interface.uuid
|string
|query
|False
a|Filter by interface.uuid

|interface.name
|string
|query
|False
a|Filter by interface.name

|interface.wwpn
|string
|query
|False
a|Filter by interface.wwpn

|initiator.wwpn
|string
|query
|False
a|Filter by initiator.wwpn

|initiator.wwnn
|string
|query
|False
a|Filter by initiator.wwnn

|initiator.aliases
|string

```

```
|query
|False
a|Filter by initiator.aliases

|initiator.port_address
|string
|query
|False
a|Filter by initiator.port_address

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|igroups.uuid
|string
|query
|False
a|Filter by igroups.uuid

|igroups.name
|string
|query
|False
a|Filter by igroups.name

|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#fc_login[fc_login]]
a|

|===

```

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```

{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "next": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroups": [
        {
          "_links": {

```

```

        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "igroup1",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
],
"initiator": {
    "aliases": [
        "alias1"
    ],
    "port_address": "5060A",
    "wwnn": "2f:a0:00:a0:98:0b:56:13",
    "wwpn": "2f:a0:00:a0:98:0b:56:13"
},
"interface": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "wwpn": "20:00:00:50:56:b4:13:a8"
},
"protocol": "string",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name

```



```

|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

```

```

[#igroups]
[.api-collapsible-fifth-title]
igroups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the initiator group.

|uuid
|string
a|The unique identifier of the initiator group.

|===

[#initiator]
[.api-collapsible-fifth-title]
initiator

Information about the logged in FC initiator.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aliases
|array[string]
a|The logged in initiator world wide port name (WWPN) aliases.

|port_address

```

```
|string
a|The port address of the initiator's FC port.
```

Each port in an FC switched fabric has its own unique port address for routing purposes. The port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the initiator port.

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a hexadecimal encoded numeric value.

```
|wwnn
|string
a|The logged in initiator world wide node name (WWNN).
```

```
|wwpn
|string
a|The logged in initiator WWPN.
```

```
|===
```

```
[#interface]
[.api-collapsible-fifth-title]
interface
```

An FC interface.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the FC interface.
```

```
|uuid
|string
a|The unique identifier of the FC interface.
```

```
|wwpn
|string
a|The WWPN of the FC interface.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#fc_login]
[.api-collapsible-fifth-title]
fc_login
```

A Fibre Channel (FC) login represents a connection formed by an FC

initiator that has successfully logged in to ONTAP. This represents the FC login on which higher-level protocols such as Fibre Channel Protocol and NVMe over Fibre Channel (NVMe/FC) rely.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|igroups
|array[link:#igroups[igroups]]
a|The initiator groups in which the initiator is a member.

|initiator
|link:#initiator[initiator]
a|Information about the logged in FC initiator.

|interface
|link:#interface[interface]
a|An FC interface.

|protocol
|string
a|The data protocol used to perform the login.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

|===

//end collapsible .Definitions block
=====

[[ID9072022d808b33f694f53152dc30fb3d]]
= Retrieve an FC login

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/fc/logins/{interface.uuid}/{initiator.wwpn}`#

Introduced In: 9.6

Retrieves an FC login.

== Related ONTAP commands

* `vserver fcp initiator show`

== Learn more

* SAN: xref:{relative_path}network_fc_logins_endpoint_overview.html[DOC
/network/fc/logins]

* NVMe: xref:{relative_path}network_fc_logins_endpoint_overview.html[DOC
/network/fc/logins]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|interface.uuid

|string

|path

|True

a|The unique identifier of the FC interface through which the initiator

logged in.

```
|initiator.wwpn
|string
|path
|True
a|The world wide port name (WWPN) of the initiator.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|igroups
|array[link:#igroups[igroups]]
a|The initiator groups in which the initiator is a member.

|initiator
|link:#initiator[initiator]
a|Information about the logged in FC initiator.

|interface
|link:#interface[interface]
a|An FC interface.
```



```
|protocol
|string
a|The data protocol used to perform the login.
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroups": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "initiator": {
    "aliases": [
      "alias1"
    ],
    "port_address": "5060A",
    "wwnn": "2f:a0:00:a0:98:0b:56:13",
    "wwpn": "2f:a0:00:a0:98:0b:56:13"
  },
  "interface": {
```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "lif1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
    "wwpn": "20:00:00:50:56:b4:13:a8"
  },
  "protocol": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

4	
---	--

	The Fibre Channel login specified does not exist.
--	---------------------------------------------------

5373983	
---------	--

	An invalid WWPN was supplied.
--	-------------------------------

5374881	
---------	--

	The Fibre Channel interface specified does not exist.
--	-------------------------------------------------------

|===

[cols=3*,options=header]

|===

Name

```

|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#igroups]
[.api-collapsible-fifth-title]

```

igroups

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the initiator group.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#initiator]
```

```
[.api-collapsible-fifth-title]
```

```
initiator
```

Information about the logged in FC initiator.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|aliases
```

```
|array[string]
```

```
a|The logged in initiator world wide port name (WWPN) aliases.
```

```
|port_address
```

```
|string
```

```
a|The port address of the initiator's FC port.
```

Each port in an FC switched fabric has its own unique port address for routing purposes. The port address is assigned by a switch in the fabric when that port logs in to the fabric. This property refers to the address given by a switch to the initiator port.

This is useful for obtaining statistics and diagnostic information from FC switches.

This is a hexadecimal encoded numeric value.

```
|wwnn
|string
a|The logged in initiator world wide node name (WWNN).
```

```
|wwpn
|string
a|The logged in initiator WWPN.
```

```
|===
```

```
[#interface]
[.api-collapsible-fifth-title]
interface
```

An FC interface.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the FC interface.
```

```
|uuid
```

```
|string
a|The unique identifier of the FC interface.
```

```
|wwpn
|string
a|The WWPN of the FC interface.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===


```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage FC WWPN aliases
```

```
:leveloffset: +1
```

```
[[ID55e709da8cc3a49e20c8bb55acdedcf3]]
```

```
= Network FC wwpn-aliases endpoint overview
```

```
== Overview
```

A worldwide port name (WWPN) is a unique 64-bit identifier for a Fibre Channel (FC) initiator. It is displayed as a 16-character hexadecimal value. SAN administrators might find it easier to identify FC initiators using an alias, especially in larger SANs.

The WWPN alias REST API allows you to create, delete, and discover aliases for WWPNS.

Multiple aliases can be created for a WWPN, but you cannot use the same alias for multiple WWPNS.

An alias can consist of up to 32 characters. Valid characters are:

- * A through Z
- * a through z
- * numbers 0 through 9
- * hyphen ("-")
- * underscore ("_")
- * left and right braces ("{" , "}")
- * period (".")

```
== Examples
```

```
=== Creating a WWPN alias
```

```
----
```

```

# The API:
POST /api/network/fc/wwpn-aliases

# The call:
curl -X POST "https://<mgmt-ip>/api/network/fc/wwpn-aliases" -H "accept:
application/json" -d '{ "svm": { "name": "svm1" }, "wwpn":
"50:0a:09:82:b4:30:25:05", "alias": "alias3" }'
-----

=== Retrieving all properties of all WWPN aliases

The `fields` query parameter is used to request that all properties be
returned.

-----

# The API:
GET /api/network/fc/wwpn-aliases

# The call:
curl -X GET "https://<mgmt-ip>/api/network/fc/wwpn-aliases?fields=*" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "68589d3d-7efa-11e8-9eed-005056b43025",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/68589d3d-7efa-11e8-9eed-005056b43025"
          }
        }
      },
      "alias": "alias1",
      "wwpn": "20:00:00:50:56:b4:30:25",
      "_links": {
        "self": {
          "href": "/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-
005056b43025/alias1"
        }
      }
    },
  ],

```

```

{
  "svm": {
    "uuid": "68589d3d-7efa-11e8-9eed-005056b43025",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/68589d3d-7efa-11e8-9eed-005056b43025"
      }
    }
  },
  "alias": "alias2",
  "wwpn": "50:0a:09:82:b4:30:25:00",
  "_links": {
    "self": {
      "href": "/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-005056b43025/alias2"
    }
  }
}
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/network/fc/wwpn-aliases"
  }
}
}
-----
'''

```

=== Retrieving all WWPN aliases named "alias1"

The `alias` query parameter is used to specify a query for the value "alias1".

The API:

GET /api/network/fc/wwpn-aliases

The call:

```

curl -X GET "https://<mgmt-ip>/api/network/fc/wwpn-aliases?alias=alias1"
-H "accept: application/hal+json"

```

The response:

```
{
"records": [
  {
    "svm": {
      "uuid": "68589d3d-7efa-11e8-9eed-005056b43025",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/68589d3d-7efa-11e8-9eed-005056b43025"
        }
      }
    },
    "alias": "alias1",
    "wwpn": "20:00:00:50:56:b4:30:25",
    "_links": {
      "self": {
        "href": "/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-005056b43025/alias1"
      }
    }
  }
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/network/fc/wwpn-aliases?alias=alias1"
  }
}
}
}
----

'''
```

=== Retrieving a specific WWPN alias

The alias to be returned is identified by the UUID of its SVM and the alias name.

The API:

GET /api/network/fc/wwpn-aliases/{svm.uuid}/{alias}

The call:

curl -X GET "https://<mgmt-ip>/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-005056b43025/alias2" -H "accept: application/hal+json"

```
# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "68589d3d-7efa-11e8-9eed-005056b43025",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/68589d3d-7efa-11e8-9eed-005056b43025"
          }
        }
      },
      "alias": "alias2",
      "wwpn": "50:0a:09:82:b4:30:25:00",
      "_links": {
        "self": {
          "href": "/api/network/fc/wwpn-aliases/68589d3d-7efa-11e8-9eed-005056b43025/alias1"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/network/fc/wwpn-aliases?alias=alias1"
    }
  }
}

```

'''

=== Deleting a WWPN alias

The alias to delete is identified by the UUID of its SVM and the alias name.

The API:

DELETE /api/network/fc/wwpn-aliases/{svm.uuid}/{alias}

The call:

curl -X DELETE "https://<mgmt-ip>/api/network/fc/wwpn-aliases/68589d3d-

```
7efa-11e8-9eed-005056b43025/alias2" -H "accept: application/hal+json"
```

```
----
```

```
[[ID6e2c7d1556f8bf56ad63ac06ba4bc0ab]]
```

```
= Retrieve FC WWPN aliases
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/network/fc/wwpn-aliases`#
```

```
*Introduced In:* 9.6
```

```
Retrieves FC WWPN aliases.
```

```
== Related ONTAP commands
```

```
* `vserver fcp wwpn-alias show`
```

```
== Learn more
```

```
* xref:{relative_path}network_fc_wwpn-aliases_endpoint_overview.html[DOC  
/network/fc/wwpn-aliases]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|alias
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by alias
```

```
|wwpn
```

```
|string
```

```
|query
|False
a|Filter by wwpn
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
```

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

- * Default value: 1
- * Max value: 120
- * Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#wwpn_alias[wwpn_alias]]
a|

|===

.Example response


```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "alias": "host1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "wwpn": "2f:a0:00:a0:98:0b:56:13"
    }
  ]
}
====

== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

```

```

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.

```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#wwpn_alias]
[.api-collapsible-fifth-title]
wwpn_alias
```

A Fibre Channel (FC) world wide port name (WWPN) alias. A WWPN is a unique 64-bit identifier for an FC initiator. It is displayed as a 16-character hexadecimal value. SAN administrators may find it easier to identify FC initiators using an alias, especially in larger SANs.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|alias
|string
a|The FC WWPN alias. Required in POST.
```

```
|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.
```

```
|wwpn
|string
a|The FC initiator WWPN. Required in POST.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[ID8ffbbbf835f2469b9173111b5f3b1582b]]
= Create an FC WWPN alias
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/network/fc/wwpn-aliases`#
```

Introduced In: 9.6

Creates an FC WWPN alias.

== Required properties

- * `svm.uuid` or `svm.name` - Existing SVM in which to create the FC alias.
- * `alias` - Name of the FC alias.
- * `wwpn` - FC WWPN for which to create the alias.

== Related ONTAP commands

```
* `vserver fcp wwpn-alias set`
```

== Learn more

```
* xref:{relative_path}network_fc_wwpn-aliases_endpoint_overview.html[DOC  
/network/fc/wwpn-aliases]
```

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|return_records
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|The default is false. If set to true, the records are returned.
```

```
* Default value:
```

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|alias
```

```
|string
```

```
a|The FC WWPN alias. Required in POST.
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|SVM, applies only to SVM-scoped objects.
```

```
|wwpn
|string
a|The FC initiator WWPN. Required in POST.
```

```
|===
```

```
.Example request
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "alias": "host1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "wwpn": "2f:a0:00:a0:98:0b:56:13"
}
```

```
====
```

```
== Response
```

Status: 201, Created


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#wwpn_alias[wwpn_alias]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "alias": "host1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "wwpn": "2f:a0:00:a0:98:0b:56:13"
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

```

```

| 1254317
| The alias already exists.

| 1260882
| The supplied SVM does not exist.

| 2621462
| The supplied SVM does not exist.

| 2621706
| Both the SVM UUID and SVM name were supplied, but do not refer to the
same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5373982
| An invalid WWPN was supplied. The valid WWN format is
XX:XX:XX:XX:XX:XX:XX, where X is a hexadecimal digit. Example:
"01:02:03:04:0a:0b:0c:0d".
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#svm]
[.api-collapsible-fifth-title]
svm
```

SVM, applies only to SVM-scoped objects.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#wwpn_alias]
[.api-collapsible-fifth-title]
wwpn_alias
```

A Fibre Channel (FC) world wide port name (WWPN) alias. A WWPN is a unique 64-bit identifier for an FC initiator. It is displayed as a 16-character hexadecimal value. SAN administrators may find it easier to identify FC initiators using an alias, especially in larger SANs.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|alias
|string
```

a|The FC WWPN alias. Required in POST.

|svm

|link:#svm[svm]

a|SVM, applies only to SVM-scoped objects.

|wwpn

|string

a|The FC initiator WWPN. Required in POST.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

```

|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

[[ID89a789ad7b2fe4c040b1e1f8f3bf18a9]]
= Delete an FC WWPN alias

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/network/fc/wwpn-aliases/{svm.uuid}/{alias}`#

*Introduced In:* 9.6

Deletes an FC WWPN alias.

== Related ONTAP commands

* `vserver fcp wwpn-alias remove`

== Learn more

* xref:{relative_path}network_fc_wwpn-aliases_endpoint_overview.html[DOC

```

```
/network/fc/wwpn-aliases]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|svm.uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|The unique identifier of the SVM.
```

```
|alias
```

```
|string
```

```
|path
```

```
|True
```

```
a|The name of FC WWPN alias.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

```
ONTAP Error Response Codes
```

```
|===
```

```
| Error Code | Description
```

```
| 1260882
```

```
| An SVM with the specified UUID does not exist.
```

```

| 5374046
| The alias could not be found.
|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]

```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message


```

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[ID4251aa0cdb25be01e3cdc517e03c5a84]]
= Retrieve an FC WWPN alias

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/network/fc/wwpn-aliases/{svm.uuid}/{alias}`#

*Introduced In:* 9.6

Retrieves an FC WWPN alias.

== Related ONTAP commands

* `vserver fcp wwpn-alias show`

== Learn more

* xref:{relative_path}network_fc_wwpn-aliases_endpoint_overview.html[DOC
/network/fc/wwpn-aliases]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path

```

```

|True
a|The unique identifier of the SVM in which the alias is found.

|alias
|string
|path
|True
a|The name of FC WWPN alias.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|alias
|string
a|The FC WWPN alias. Required in POST.

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.

|wwpn
|string
a|The FC initiator WWPN. Required in POST.

```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "alias": "host1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "wwpn": "2f:a0:00:a0:98:0b:56:13"
}
```

```
====
```

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 1260882
```

```
| The supplied SVM does not exist.
```

```
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid

```

```
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage FC services for SVMs
```

```
:leveloffset: +1
```

```
[[ID239e9a75ade8c8b22072a15d57185469]]
```

```
= Protocols SAN fcp services endpoint overview
```

```
== Overview
```

A Fibre Channel Protocol (FC Protocol) service defines the properties of the FC Protocol target for an SVM. There can be at most one FC Protocol service for an SVM. An SVM FC Protocol service must be created before FC Protocol initiators can log in to the SVM.

The FC Protocol service REST API allows you to create, update, delete, and discover FC services for SVMs.

```
== Performance monitoring
```

Performance of the SVM can be monitored by the `metric.+++` and

`statistics.+++` properties. These show the performance of the SVM in terms of IOPS, latency, and throughput. The `metric.+++` properties denote an average whereas `statistics.+++` properties denote a real-time monotonically increasing value aggregated across all nodes.

== Examples

=== Creating an FC Protocol service for an SVM

The simplest way to create an FC Protocol service is to specify only the SVM, either by name or UUID. By default, the new FC Protocol service is enabled.

In this example, the `return_records` query parameter is used to retrieve the new FC Protocol service object in the REST response.

The API:

POST /api/protocols/san/fcp/services

The call:

```
curl -X POST 'https://<mgmt-  
ip>/api/protocols/san/fcp/services?return_records=true' -H 'accept:  
application/hal+json' -d '{ "svm": { "name": "svm1" } }'
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "svm": {  
        "uuid": "5c659d90-c01a-11e8-88ed-005056bbb24b",  
        "name": "svm1",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/5c659d90-c01a-11e8-88ed-005056bbb24b"  
          }  
        }  
      },  
      "enabled": true,  
      "target": {  
        "name": "20:00:00:50:56:bb:b2:4b"  
      },  
      "_links": {  
        "self": {
```



```

        "href": "/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-005056bbb24b"
    }
}
]
}
----

'''

=== Retrieving FC Protocol services for all SVMs in the cluster

----

# The API:
GET /api/protocols/san/fcp/services

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/fcp/services' -H 'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "5c659d90-c01a-11e8-88ed-005056bbb24b",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/5c659d90-c01a-11e8-88ed-005056bbb24b"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-005056bbb24b"
        }
      }
    },
    {
      "svm": {
        "uuid": "6011f874-c01a-11e8-88ed-005056bbb24b",
        "name": "svm2",

```

```

    "_links": {
      "self": {
        "href": "/api/svm/svms/6011f874-c01a-11e8-88ed-005056bbb24b"
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/fcp/services/6011f874-c01a-11e8-88ed-005056bbb24b"
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/protocols/san/fcp/services"
    }
  }
}
----

'''

```

=== Retrieving details for a specific FC Protocol service

The FC Protocol service is identified by the UUID of its SVM.

The API:

GET /api/protocols/san/fcp/services/{svm.uuid}

The call:

curl -X GET 'https://<mgmt-ip>/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-005056bbb24b' -H 'accept: application/hal+json'

The response:

```

{
  "svm": {
    "uuid": "5c659d90-c01a-11e8-88ed-005056bbb24b",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5c659d90-c01a-11e8-88ed-005056bbb24b"
      }
    }
  }
}

```

```

    }
  }
},
"enabled": true,
"target": {
  "name": "20:00:00:50:56:bb:b2:4b"
},
"_links": {
  "self": {
    "href": "/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-005056bbb24b"
  }
}
}
}

```

'''

=== Disabling an FC Protocol service

Disabling an FC Protocol service shuts down all active FC Protocol logins for the SVM and prevents new FC Protocol logins.

The FC Protocol service to update is identified by the UUID of its SVM.

The API:

PATCH /api/protocols/san/fcp/services/{svm.uuid}

The call:

```

curl -X PATCH 'https://<mgmt-ip>/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-005056bbb24b' -H 'accept: application/hal+json' -d '{"enabled": "false"}'

```

You can retrieve the FC Protocol service to confirm the change.

In this example, the `fields` query parameter is used to limit the response to the `enabled` property and FC Protocol service identifiers.

The API:

GET /api/protocols/san/fcp/services/{svm.uuid}

The call:

```
curl -X GET 'https://<mgmt-ip>/api/protocols/san/fcp/services/5c659d90-
c01a-11e8-88ed-005056bbb24b?fields=enabled' -H 'accept:
application/hal+json'
```

The response:

```
{
  "svm": {
    "uuid": "5c659d90-c01a-11e8-88ed-005056bbb24b",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5c659d90-c01a-11e8-88ed-005056bbb24b"
      }
    }
  },
  "enabled": false,
  "_links": {
    "self": {
      "href": "/api/protocols/san/fcp/services/5c659d90-c01a-11e8-88ed-
005056bbb24b"
    }
  }
}
```

'''

=== Deleting an FC Protocol service

The FC Protocol service must be disabled before it can be deleted.

The FC Protocol service to delete is identified by the UUID of its SVM.

The API:

```
DELETE /api/protocols/san/fcp/services/{svm.uuid}
```

The call:

```
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/fcp/services/5c659d90-
c01a-11e8-88ed-005056bbb24b' -H 'accept: application/hal+json'
```

```
[[IDa55d2f043943c5726213a2b9792ba510]]
```

= Retrieve FC protocol services

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/protocols/san/fcp/services`#
```

Introduced In: 9.6

Retrieves FC Protocol services.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `statistics.+++`

* `metric.+++`

== Related ONTAP commands

* `vserver fcp show`

== Learn more

*

[xref:{relative_path}protocols_san_fcp_services_endpoint_overview.html\[DOC /protocols/san/fcp/services\]](#)

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|svm.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write

* Introduced in: 9.7

|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read

* Introduced in: 9.7

|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total

* Introduced in: 9.7

|metric.duration
|string
|query
|False
a|Filter by metric.duration

* Introduced in: 9.7

|metric.timestamp
```

```
|string
|query
|False
a|Filter by metric.timestamp
```

* Introduced in: 9.7

```
|metric.status
|string
|query
|False
a|Filter by metric.status
```

* Introduced in: 9.7

```
|metric.iops.total
|integer
|query
|False
a|Filter by metric.iops.total
```

* Introduced in: 9.7

```
|metric.iops.read
|integer
|query
|False
a|Filter by metric.iops.read
```

* Introduced in: 9.7

```
|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other
```

* Introduced in: 9.7

```
|metric.iops.write
|integer
|query
```

```
|False
a|Filter by metric.iops.write

* Introduced in: 9.7

|metric.latency.total
|integer
|query
|False
a|Filter by metric.latency.total

* Introduced in: 9.7

|metric.latency.read
|integer
|query
|False
a|Filter by metric.latency.read

* Introduced in: 9.7

|metric.latency.other
|integer
|query
|False
a|Filter by metric.latency.other

* Introduced in: 9.7

|metric.latency.write
|integer
|query
|False
a|Filter by metric.latency.write

* Introduced in: 9.7

|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total
```


* Introduced in: 9.7

```
|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read
```

* Introduced in: 9.7

```
|statistics.latency_raw.other
|integer
|query
|False
a|Filter by statistics.latency_raw.other
```

* Introduced in: 9.7

```
|statistics.latency_raw.write
|integer
|query
|False
a|Filter by statistics.latency_raw.write
```

* Introduced in: 9.7

```
|statistics.timestamp
|string
|query
|False
a|Filter by statistics.timestamp
```

* Introduced in: 9.7

```
|statistics.iops_raw.total
|integer
|query
|False
a|Filter by statistics.iops_raw.total
```

* Introduced in: 9.7

```
|statistics.iops_raw.read
|integer
|query
|False
a|Filter by statistics.iops_raw.read
```

* Introduced in: 9.7

```
|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other
```

* Introduced in: 9.7

```
|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write
```

* Introduced in: 9.7

```
|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write
```

* Introduced in: 9.7

```
|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read
```

* Introduced in: 9.7

```
|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total
```

* Introduced in: 9.7

```
|statistics.status
|string
|query
|False
a|Filter by statistics.status
```

* Introduced in: 9.7

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|target.name
|string
|query
|False
a|Filter by target.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
```

```

|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|num_records
|integer
a|Number of records.

|records
|array[link:#fcp_service[fcp_service]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "duration": "PT15S",
        "iops": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "latency": {
          "read": "200",

```

```

        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"statistics": {
    "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"target": {
    "name": "20:00:00:50:56:bb:b2:4b"
}
}
]
}
=====

```

```
== Error
```

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
```

```

href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self

```



```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#iops]
```

```
[.api-collapsible-fifth-title]
```

```
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput]
```

```
[.api-collapsible-fifth-title]
```

throughput

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```

|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be
```

metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
```

[.api-collapsible-fifth-title]

statistics

[cols=3*,options=header]

|===

|Name

|Type

|Description

|iops_raw

|link:#iops_raw[iops_raw]

a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw

|link:#latency_raw[latency_raw]

a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#target]

[.api-collapsible-fifth-title]

target

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The target name of the FC Protocol service. This is generated for the SVM during POST.

The FC Protocol target name is a world wide node name (WWNN).

If required, the target name can be modified using the ONTAP command line.

- * example: 20:00:00:50:56:bb:b2:4b
- * maxLength: 128
- * minLength: 1
- * readOnly: 1
- * Introduced in: 9.6

|===

```
[#fcp_service]
[.api-collapsible-fifth-title]
fcp_service
```

A Fibre Channel (FC) Protocol service defines the properties of the FC Protocol target for an SVM. There can be at most one FC Protocol service for an SVM. An SVM's FC Protocol service must be created before FC Protocol initiators can login to the SVM.

A FC Protocol service is identified by the UUID of its SVM.

```
[cols=3*,options=header]
```

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|enabled

|boolean

a|The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

|metric

```

|link:#metric[metric]
a|

|statistics
|link:#statistics[statistics]
a|

|svm
|link:#svm[svm]
a|

|target
|link:#target[target]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===

```

```

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

[[ID24ffdfa4ead03f0ef83bad338c121915]]
= Create an FC protocol service

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/san/fcp/services`#

*Introduced In:* 9.6

Creates an FC Protocol service.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the FC
Protocol service.

```

== Related ONTAP commands

* `vserver fcp create`

== Learn more

*

xref:{relative_path}protocols_san_fcp_services_endpoint_overview.html[DOC
/protocols/san/fcp/services]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
|link:#metric[metric]
a|

|statistics
|link:#statistics[statistics]
a|

|svm
|link:#svm[svm]
a|

|target
|link:#target[target]
a|

|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "metric": {
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
  },
}
```

```

    "timestamp": "2017-01-25T11:20:13Z"
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "name": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "name": "20:00:00:50:56:bb:b2:4b"
  }
}
====

== Response

```

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

```

```

|records
|array[link:#fcp_service[fcp_service]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "metric": {
        "duration": "PT15S",
        "iops": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "latency": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "status": "ok",
        "throughput": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "timestamp": "2017-01-25T11:20:13Z"
      },
      "statistics": {
        "iops_raw": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "latency_raw": {
          "read": "200",
          "total": "1000",
          "write": "100"
        },
        "status": "ok",

```

```

    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "name": "20:00:00:50:56:bb:b2:4b"
  }
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
|
| 1115127
| The cluster lacks a valid FCP license.
|
| 2621462
| The supplied SVM does not exist.
|
| 2621507
| The Fibre Channel Protocol is not allowed for the specified SVM.
|
| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.
|
| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.
|
| 5374082
| The Fibre Channel Protocol service already exists for the SVM.

```



```
| 5374092
| The Fibre Channel Procotol is not supported on the cluster hardware
configuration; there are no Fibre Channel adapters.

| 5374893
| The SVM is stopped. The SVM must be running to create a Fibre Channel
Protocol service.
|===
```

```
== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
```

```

|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read

```

```
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric

[cols=3*,options=header]
|===
|Name
|Type
|Description

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#statistics]
[.api-collapsible-fifth-title]
statistics

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency

```

per I/O operation.

|status

|string

a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.


```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#target]
[.api-collapsible-fifth-title]
target
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The target name of the FC Protocol service. This is generated for the
SVM during POST.
```

The FC Protocol target name is a world wide node name (WWNN).

If required, the target name can be modified using the ONTAP command line.

```
* example: 20:00:00:50:56:bb:b2:4b
* maxLength: 128
* minLength: 1
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#fcp_service]
[.api-collapsible-fifth-title]
fcp_service
```

A Fibre Channel (FC) Protocol service defines the properties of the FC Protocol target for an SVM. There can be at most one FC Protocol service for an SVM. An SVM's FC Protocol service must be created before FC Protocol initiators can login to the SVM.

A FC Protocol service is identified by the UUID of its SVM.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

a|The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
```

```
|link:#metric[metric]
```

```
a|
```

```
|statistics
```

```
|link:#statistics[statistics]
```

```
a|
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|target
```

```
|link:#target[target]
```

```
a|
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

```
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

```

====

[[ID99cb2393e711cefb0647f8b40adbd3a7]]

= Delete an FC protocol service

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-block]#`/protocols/san/fcp/services/{svm.uuid}`#

Introduced In: 9.6

Deletes an FC Protocol service. An FC Protocol service must be disabled before it can be deleted.

== Related ONTAP commands

* `vserver fcp delete`

== Learn more

*

xref:{relative_path}protocols_san_fcp_services_endpoint_overview.html[DOC /protocols/san/fcp/services]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|The unique identifier of the SVM for which to delete the FC Protocol service.

|==

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

2621462	
---------	--

An SVM with the specified UUID does not exist.	
------------------------------------------------	--

5373960	
---------	--

The Fibre Channel Protocol service cannot be removed while it is enabled.	
---------------------------------------------------------------------------	--

5374083	
---------	--

There is no Fibre Channel Protocol service for the specified SVM.	
-------------------------------------------------------------------	--

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "error": {

 "arguments": [

 {

```

        "code": "string",
        "message": "string"
    }
],
"code": "4",
"message": "entry doesn't exist",
"target": "uuid"
}
}
=====

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====

```

```

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

```

```

|===

```

```

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[IDfc56a8cbd2d6f9c766e496460ddb22c7]]
= Retrieve an FC protocol service
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/san/fcp/services/{svm.uuid}`#
```

```
*Introduced In:* 9.6
```

```
Retrieves an FC Protocol service.
```

```
== Related ONTAP commands
```

```
* `vserver fcp show`
```

```
== Learn more
```

*

xref:{relative_path}protocols_san_fcp_services_endpoint_overview.html[DOC
/protocols/san/fcp/services]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|The unique identifier of the SVM for which to retrieve the FC Protocol service.

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|


```
|enabled
|boolean
a|The administrative state of the FC Protocol service. The FC Protocol
service can be disabled to block all FC Protocol connectivity to the SVM.
```

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
|link:#metric[metric]
a|
```

```
|statistics
|link:#statistics[statistics]
a|
```

```
|svm
|link:#svm[svm]
a|
```

```
|target
|link:#target[target]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "duration": "PT15S",
```

```

    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },

```

```

    "target": {
      "name": "20:00:00:50:56:bb:b2:4b"
    }
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 2621462
| An SVM with the specified UUID does not exist.
|
| 5374083
| There is no Fibre Channel Protocol service for the specified SVM.
|====

```

```
[cols=3*,options=header]
```

```

|====
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```
|====
```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",

```

```

        "message": "string"
    }
],
"code": "4",
"message": "entry doesn't exist",
"target": "uuid"
}
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]

```

a|

|===

[#iops]

[.api-collapsible-fifth-title]

iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#latency]

[.api-collapsible-fifth-title]

latency

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput]
```

```
[.api-collapsible-fifth-title]
```

throughput

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#metric]

[.api-collapsible-fifth-title]

metric

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|duration

|string

a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

|iops

|link:#iops[iops]

a|The rate of I/O operations observed at the storage object.

|latency

|link:#latency[latency]

a|The round trip latency in microseconds observed at the storage object.

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.


```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
```

```

statistics

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.

|timestamp
|string
a|The timestamp of the performance data.

```

```
|===
```

```
[#svm]  
[.api-collapsible-fifth-title]  
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|name  
|string  
a|The name of the SVM.
```

```
|uuid  
|string  
a|The unique identifier of the SVM.
```

```
|===
```

```
[#target]  
[.api-collapsible-fifth-title]  
target
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|name  
|string  
a|The target name of the FC Protocol service. This is generated for the  
SVM during POST.
```

The FC Protocol target name is a world wide node name (WWNN).

If required, the target name can be modified using the ONTAP command line.

- * example: 20:00:00:50:56:bb:b2:4b
- * maxLength: 128
- * minLength: 1
- * readOnly: 1
- * Introduced in: 9.6

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

|===

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID75c73917a4ccb4869c0d515437388b43]]
= Update an FC protocol service

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/san/fcp/services/{svm.uuid}`#

*Introduced In:* 9.6

Updates an FC Protocol service.

== Related ONTAP commands

* `vserver fcp modify`
* `vserver fcp start`
* `vserver fcp stop`

== Learn more

*

```

xref:{relative_path}protocols_san_fcp_services_endpoint_overview.html[DOC
/protocols/san/fcp/services]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True

a|The unique identifier of the SVM whose FC Protocol service is to be updated.

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean

a|The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

|metric
|link:#metric[metric]
a|

```
|statistics
|link:#statistics[statistics]
a|
```

```
|svm
|link:#svm[svm]
a|
```

```
|target
|link:#target[target]
a|
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
```

```
  "metric": {
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
  },
```



```

    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "name": "20:00:00:50:56:bb:b2:4b"
  }
}
====

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 2621462
| An SVM with the specified UUID does not exist.
|
| 5374083
| There is no Fibre Channel Protocol service for the specified SVM.
|====

```

```

== Definitions

```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be  
metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|read
|integer

a|Performance metric for read I/O operations.

|total
|integer

a|Performance metric aggregated over all types of I/O operations.

|write
|integer

a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

[cols=3*,options=header]

|===

|Name
|Type
|Description

```

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#iops_raw]
[.api-collapsible-fifth-title]

```

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
```

```
[.api-collapsible-fifth-title]
```

latency_raw

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This should be used along
with delta time to calculate the rate of throughput bytes per unit of
time.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.

```

```

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#statistics]
[.api-collapsible-fifth-title]
statistics

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time

```


between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#target]
[.api-collapsible-fifth-title]
```

target

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The target name of the FC Protocol service. This is generated for the SVM during POST.

The FC Protocol target name is a world wide node name (WWNN).

If required, the target name can be modified using the ONTAP command line.

* example: 20:00:00:50:56:bb:b2:4b

* maxLength: 128

* minLength: 1

* readOnly: 1

* Introduced in: 9.6

|===

[#fcp_service]

[.api-collapsible-fifth-title]

fcp_service

A Fibre Channel (FC) Protocol service defines the properties of the FC Protocol target for an SVM. There can be at most one FC Protocol service for an SVM. An SVM's FC Protocol service must be created before FC Protocol initiators can login to the SVM.

A FC Protocol service is identified by the UUID of its SVM.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|The administrative state of the FC Protocol service. The FC Protocol service can be disabled to block all FC Protocol connectivity to the SVM.

This is optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
|link:#metric[metric]
a|
```

```
|statistics
|link:#statistics[statistics]
a|
```

```
|svm
|link:#svm[svm]
a|
```

```
|target
|link:#target[target]
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDbd97e21e2fe2a0db2426a121c89e47cc]]
= Retrieve historical performance metrics for the FC protocol service of
an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/san/fcp/services/{svm.uuid}/metrics`#

```

Introduced In: 9.7

Retrieves historical performance metrics for the FC Protocol service of an SVM.

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|timestamp

|string

|query

|False

a|Filter by timestamp

|duration

|string

|query

|False

a|Filter by duration

|throughput.total

|integer

|query

|False

a|Filter by throughput.total

|throughput.read

|integer

|query

|False

a|Filter by throughput.read

|throughput.other

|integer

```
|query  
|False  
a|Filter by throughput.other
```

```
|throughput.write  
|integer  
|query  
|False  
a|Filter by throughput.write
```

```
|latency.total  
|integer  
|query  
|False  
a|Filter by latency.total
```

```
|latency.read  
|integer  
|query  
|False  
a|Filter by latency.read
```

```
|latency.other  
|integer  
|query  
|False  
a|Filter by latency.other
```

```
|latency.write  
|integer  
|query  
|False  
a|Filter by latency.write
```

```
|status  
|string  
|query  
|False  
a|Filter by status
```

```
|iops.total
|integer
|query
|False
a|Filter by iops.total
```

```
|iops.read
|integer
|query
|False
a|Filter by iops.read
```

```
|iops.other
|integer
|query
|False
a|Filter by iops.other
```

```
|iops.write
|integer
|query
|False
a|Filter by iops.write
```

```
|svm.uuid
|string
|path
|True
a|The unique identifier of the SVM.
```

```
|interval
|string
|query
|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:
```

- * 1h: Metrics over the most recent hour sampled over 15 seconds.
- * 1d: Metrics over the most recent day sampled over 5 minutes.
- * 1w: Metrics over the most recent week sampled over 30 minutes.
- * 1m: Metrics over the most recent month sampled over 2 hours.
- * 1y: Metrics over the most recent year sampled over a day.

```

* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

```

```

* Default value: 1
* Max value: 120
* Min value: 0

```

```

|fields
|array[string]
|query
|False
a|Specify the fields to return.

```

```

|max_records
|integer
|query
|False
a|Limit the number of records returned.

```

```

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

```

```

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

```


* Default value: 1

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of records

|records

|array[link:#records[records]]

a|

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "_links": {

 "next": {

 "href": "/api/resourcelink"

 },

 "self": {

 "href": "/api/resourcelink"

 }

 },

 "records": [

 {

 "_links": {

```

        "self": {
            "href": "/api/resourcelink"
        },
    },
    "duration": "PT15S",
    "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error

```

```

[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#iops]
[.api-collapsible-fifth-title]
iops

```

The rate of I/O operations observed at the storage object.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#records]  
[.api-collapsible-fifth-title]  
records
```

Performance numbers, such as IOPS latency and throughput.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|duration  
|string  
a|The duration over which this sample is calculated. The time durations  
are represented in the ISO-8601 standard format. Samples can be calculated  
over the following durations:
```

```
|iops  
|link:#iops[iops]  
a|The rate of I/O operations observed at the storage object.
```

```
|latency  
|link:#latency[latency]  
a|The round trip latency in microseconds observed at the storage object.
```

```
|status  
|string  
a|Errors associated with the sample. For example, if the aggregation of  
data over multiple nodes fails, then any partial errors might return "ok"  
on success or "error" on an internal uncategorized failure. Whenever a  
sample collection is missed but done at a later time, it is back filled to  
the previous 15 second timestamp and tagged with "backfilled_data".  
"Inconsistent_delta_time" is encountered when the time between two  
collections is not the same for all nodes. Therefore, the aggregated value
```

might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====

:leveloffset: -1

= Manage SAN igroups

:leveloffset: +1

[[ID26f0f42530d9df890f2a43ea11dd3ee8]]
= Protocols SAN igroups endpoint overview


== Overview

```

An initiator group (igroup) is a collection of Fibre Channel (FC) world wide port names (WWPNs), and/or iSCSI Qualified Names (IQNs), and/or iSCSI EUIs (Extended Unique Identifiers) that identify host initiators.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, create an initiator group containing the host initiator names, then create a LUN map that associates the initiator group with the LUN.

The initiator group REST API allows you to create, update, delete, and discover initiator groups, and add and remove initiators that can access the target and associated LUNs.

An initiator can appear in multiple initiator groups. An initiator group can be mapped to multiple LUNs. A specific initiator can be mapped to a specific LUN only once.

All initiators in an initiator group must be from the same operating system. The initiator group's operating system is specified when the initiator group is created.

When an initiator group is created, the ``protocol`` property is used to restrict member initiators to Fibre Channel (`_fcp_`), iSCSI (`_iscsi_`), or both (`_mixed_`).

Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the ``/protocols/san/igroups/{igroup.uuid}/initiators`` endpoint. See [xref:{relative_path}post-protocols-san-igroups-initiators\(#-san-igroup-initiator-create\)anddelete-protocols-san-igroups-initiators-.html](#)<</SAN/igroup_initiator_create,POST `/protocols/san/igroups/{igroup.uuid}/initiators>>` and `[DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}]` for more details.

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

== Examples

=== Creating an initiator group with no initiators

The example initiator group is for Linux iSCSI initiators only. Note that

the `return_records` query parameter is used to obtain the newly created initiator group in the response.

The API:

POST /api/protocols/san/igroups

The call:

```
curl -X POST 'https://<mgmt-  
ip>/api/protocols/san/igroups?return_records=true' -H 'accept:  
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name": "igroup1",  
"os_type": "linux", "protocol": "iscsi" }'
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "svm": {  
        "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",  
        "name": "svm1",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"  
          }  
        }  
      },  
      "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",  
      "name": "igroup1",  
      "protocol": "iscsi",  
      "os_type": "linux",  
      "_links": {  
        "self": {  
          "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-  
005056bb7072"  
        }  
      }  
    }  
  ]  
}
```

'''

=== Creating an initiator group with initiators

The example initiator group is for Windows. FC Protocol and iSCSI initiators are allowed. Note that the `return_records` query parameter is used to obtain the newly created initiator group in the response.

The API:

POST /api/protocols/san/igroups

The call:

```
curl -X POST 'https://<mgmt-  
ip>/api/protocols/san/igroups?return_records=true' -H 'accept:  
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name": "igroup2",  
"os_type": "windows", "protocol": "mixed", "initiators": [ { "name":  
"20:01:00:50:56:bb:70:72" }, { "name": "iqn.1991-05.com.ms:host1" } ] }'
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "svm": {  
        "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",  
        "name": "svm1",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"  
          }  
        }  
      },  
      "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",  
      "name": "igroup2",  
      "protocol": "mixed",  
      "os_type": "windows",  
      "initiators": [  
        {  
          "name": "20:01:00:50:56:bb:70:72",  
          "_links": {  
            "self": {  
              "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-  
005056bb7072/initiators/20:01:00:50:56:bb:70:72"  
            }  
          }  
        },  
        {  
          "name": "iqn.1991-05.com.ms:host1",
```

```

    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072/initiators/iqn.1991-05.com.ms:host1"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072"
    }
  }
}
]
}
----
'''

```

=== Retrieving all initiator groups

The API:

GET /api/protocols/san/igroups

The call:

curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups' -H 'accept: application/hal+json'

The response:

```

{
  "records": [
    {
      "svm": {
        "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
          }
        }
      },
      "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",

```

```

    "name": "igroup1",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072"
      }
    }
  },
  {
    "svm": {
      "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
        }
      }
    },
    "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
    "name": "igroup2",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups"
  }
}
}
----

'''

```

=== Retrieving all properties of all initiator groups

The `fields` query parameter is used to request all initiator group properties.

```

# The API:
GET /api/protocols/san/igroups

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups?fields=*' -H
'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
          }
        }
      },
      "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
      "name": "igroup1",
      "protocol": "iscsi",
      "os_type": "linux",
      "_links": {
        "self": {
          "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072"
        }
      }
    },
    {
      "svm": {
        "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
          }
        }
      },
      "uuid": "abf9c39d-ab9f-11e8-b8a3-005056bb7072",
      "name": "igroup2",
      "protocol": "mixed",
      "os_type": "windows",
      "initiators": [

```

```

{
  "name": "20:01:00:50:56:bb:70:72",
  "_links": {
    "self": {
      "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072/initiators/20:01:00:50:56:bb:70:72"
    }
  }
},
{
  "name": "iqn.1991-05.com.ms:host1",
  "_links": {
    "self": {
      "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072/initiators/iqn.1991-05.com.ms:host1"
    }
  }
}
],
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups/abf9c39d-ab9f-11e8-b8a3-005056bb7072"
  }
}
}
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups?fields=*"
  }
}
}
}

```

=== Retrieving all initiator groups for Linux

The `os_type` query parameter is used to perform the query.

```

----

# The API:
GET /api/protocols/san/igroups

```



```

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups?os_type=linux' -H
'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
          }
        }
      },
      "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
      "name": "igroup1",
      "os_type": "linux",
      "_links": {
        "self": {
          "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-
005056bb7072"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/protocols/san/igroups?os_type=linux"
    }
  }
}
----

'''

=== Retrieving a specific initiator group

----

# The API:
GET /api/protocols/san/igroups/{uuid}

```

```

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072' -H 'accept: application/hal+json'

# The response:
{
  "svm": {
    "uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
      }
    }
  },
  "uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
  "name": "igroup1",
  "protocol": "iscsi",
  "os_type": "linux",
  "_links": {
    "self": {
      "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072"
    }
  }
}
----

'''

=== Retrieving LUNs mapped to a specific initiator group

The `fields` parameter is used to specify the desired properties.

----

# The API:
GET /api/protocols/san/igroups

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072?fields=lun_maps' -H 'accept: application/hal+json'

# The response:
{
  "svm": {

```

```

"uuid": "02b0dfff-aa28-11e8-a653-005056bb7072",
"name": "svm1",
"_links": {
  "self": {
    "href": "/api/svm/svms/02b0dfff-aa28-11e8-a653-005056bb7072"
  }
},
"uuid": "8f249e7d-ab9f-11e8-b8a3-005056bb7072",
"name": "igroup1",
"lun_maps": [
  {
    "logical_unit_number": 0,
    "lun": {
      "name": "/vol/vol1/lun1",
      "uuid": "4b33ba57-c4e0-4dbb-bc47-214800d18a71",
      "node": {
        "name": "node1",
        "uuid": "f17182af-223f-4d51-8197-2cb2146d5c4c",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/f17182af-223f-4d51-8197-2cb2146d5c4c"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/storage/luns/4b33ba57-c4e0-4dbb-bc47-214800d18a71"
        }
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072"
  }
}
}
}
----

'''

=== Renaming an initiator group

```

Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.

The API:

PATCH /api/protocols/san/igroups/{uuid}

The call:

```
curl -X PATCH 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072' -H 'accept: application/hal+json' -d '{ "name": "igroup1_newName" }'
```

'''

=== Changing the operating system type of an initiator group

The API:

PATCH /api/protocols/san/igroups/{uuid}

The call:

```
curl -X PATCH 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072' -H 'accept: application/hal+json' -d '{ "os_type": "aix" }'
```

'''

=== Adding an initiator to an initiator group

The API:

POST /api/protocols/san/igroups/{igroup.uuid}/initiators

The call:

```
curl -X POST 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072/initiators' -H 'accept: application/hal+json' -d '{ "name": "iqn.1991-05.com.ms:host2" }'
```

'''

=== Adding multiple initiators to an initiator group

Note the use of the `records` property to add multiple initiators to the initiator group in a single API call.

The API:

POST /api/protocols/san/igroups/{igroup.uuid}/initiators

The call:

```
curl -X POST 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072/initiators' -H 'accept: application/hal+json' -d '{
  "records": [ { "name": "iqn.1991-05.com.ms:host3" }, { "name": "iqn.1991-05.com.ms:host4" } ] }'
```

'''

=== Removing an initiator from an initiator group

The API:

DELETE /api/protocols/san/igroups/{igroup.uuid}/initiators/iqn.1991-05.com.ms:host3

The call:

```
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-11e8-b8a3-005056bb7072/initiators/iqn.1991-05.com.ms:host3' -H 'accept: application/hal+json'
```

'''

=== Removing an initiator from a mapped initiator group

Normally, removing an initiator from an initiator group that is mapped to a LUN is not allowed. The removal can be forced using the `allow_delete_while_mapped` query parameter.

The API:

DELETE /api/protocols/san/igroups/{igroup.uuid}/initiators/iqn.1991-05.com.ms:host4

```

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/8f249e7d-ab9f-
11e8-b8a3-005056bb7072/initiators/iqn.1991-
05.com.ms:host4?allow_delete_while_mapped=true' -H 'accept:
application/hal+json'
----

'''

=== Deleting an initiator group

----

# The API:
DELETE /api/protocols/san/igroups/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/abf9c39d-ab9f-
11e8-b8a3-005056bb7072' -H 'accept: application/hal+json'
----

'''

=== Deleting a mapped initiator group

Normally, deleting an initiator group that is mapped to a LUN is not
allowed. The deletion can be forced using the `allow_delete_while_mapped`
query parameter.

----

# The API:
DELETE /api/protocols/san/igroups/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/igroups/abf9c39d-ab9f-
11e8-b8a3-005056bb7072?allow_delete_while_mapped=true' -H 'accept:
application/hal+json'
----

[[ID0cf4c7fcfbefb10ad4450198e0993fe]]
= Retrieve initiator groups

```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/san/igroups`#
```

Introduced In: 9.6

Retrieves initiator groups.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `lun_maps.+++`

== Related ONTAP commands

* `lun igroup show`

* `lun mapping show`

== Learn more

* [xref:{relative_path}protocols_san_igroups_endpoint_overview.html\[DOC /protocols/san/igroups\]](#)

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|protocol
```

```
|string
|query
|False
a|Filter by protocol
```

```
|os_type
|string
|query
|False
a|Filter by os_type
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|initiators.name
|string
|query
|False
a|Filter by initiators.name
```

```
|initiators.igroup.uuid
|string
|query
|False
a|Filter by initiators.igroup.uuid
```

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
|query
|False
a|Filter by svm.name
```



```
|delete_on_unmap
|boolean
|query
|False
a|Filter by delete_on_unmap
```

```
|lun_maps.lun.node.name
|string
|query
|False
a|Filter by lun_maps.lun.node.name
```

```
|lun_maps.lun.node.uuid
|string
|query
|False
a|Filter by lun_maps.lun.node.uuid
```

```
|lun_maps.lun.name
|string
|query
|False
a|Filter by lun_maps.lun.name
```

```
|lun_maps.lun.uuid
|string
|query
|False
a|Filter by lun_maps.lun.uuid
```

```
|lun_maps.logical_unit_number
|integer
|query
|False
a|Filter by lun_maps.logical_unit_number
```

```
|fields
|array[string]
|query
|False
```

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|return_records

|boolean

|query

|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|num_records
```

```
|integer
```

```
a|Number of records.
```

```
|records
```

```
|array[link:#igroup[igroup]]
```

```
a|
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "_links": {
```

```
    "next": {
```

```
      "href": "/api/resourcelink"
```

```
    },
```

```
    "self": {
```

```
      "href": "/api/resourcelink"
```

```
    }
```

```
  },
```

```
  "records": [
```

```
    {
```

```
      "_links": {
```

```
        "self": {
```

```
          "href": "/api/resourcelink"
```

```
        }
```

```
      },
```

```
      "initiators": [
```

```
        {
```

```
          "_links": {
```

```
            "self": {
```

```

        "href": "/api/resourcelink"
    }
},
"igroup": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
},
"name": "iqn.1998-01.com.corp.iscsi:name1"
}
],
"lun_maps": [
    {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "logical_unit_number": 0,
        "lun": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "lun1",
            "node": {
                "_links": {
                    "self": {
                        "href": "/api/resourcelink"
                    }
                },
                "name": "node1",
                "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        }
    }
],
"name": "igroup1",
"os_type": "string",
"protocol": "string",
"svm": {

```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",

```

```

    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|self  
|link:#href[href]  
a|
```

```
|===
```

```
[#igroup]  
[.api-collapsible-fifth-title]  
igroup
```

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|uuid  
|string  
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#igroup_initiator_no_records]
[.api-collapsible-fifth-title]
igroup_initiator_no_records
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|igroup
```

```
|link:#igroup[igroup]
```

```
a|The initiator group in which the initiator is found.
```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
|name
```

```
|string
```

```
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.
```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

```
|===
```

```
[#node]
```

```
[.api-collapsible-fifth-title]
```

```
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#lun]
[.api-collapsible-fifth-title]
lun
```

The LUN to which the initiator group is mapped.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the LUN.
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|The unique identifier of the LUN.
```

```
|===
```

```
[#lun_maps]
[.api-collapsible-fifth-title]
lun_maps
```

A LUN map with which the initiator group is associated.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|logical_unit_number
|integer
a|The logical unit number assigned to the LUN for initiators in the
initiator group.
```

```
|lun
|link:#lun[lun]
a|The LUN to which the initiator group is mapped.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

An initiator group (igroup) is a collection of Fibre Channel (FC) world wide port names (WWPN), and/or iSCSI Qualified Names (IQNs), and/or iSCSI EUIs (Extended Unique Identifiers) that identify host initiators.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, create an initiator group containing the hosts' initiator names, then create a LUN map that associates the initiator group with the LUN.

An initiator can appear in multiple initiator groups. An initiator group can be mapped to multiple LUNs. A specific initiator can be mapped to a specific LUN only once.

All initiators in an initiator group must be from the same operating system. The initiator group's operating system is specified when the initiator group is created.

When an initiator group is created, the `protocol` property is used to restrict member initiators to Fibre Channel (`_fcp_`), iSCSI (`_iscsi_`), or both (`_mixed_`).

Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the ``/protocols/san/igroups/{igroup.uuid}/initiators`` endpoint. See [xref:{relative_path}post-protocols-san-igroups-initiators\(#-san-igroup-initiator-create\)anddelete-protocols-san-igroups-initiators-.html](#)<</SAN/igroup_initiator_create,POST `/protocols/san/igroups/{igroup.uuid}/initiators``>> and `[DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}]` for more

details.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|delete_on_unmap
```

```
|boolean
```

a|An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to `_false_` when the initiator group is created.

```
|initiators
```

```
|array[link:#igroup_initiator_no_records[igroup_initiator_no_records]]
```

a|The initiators that are members of the group. Optional in POST.

Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the

``/protocols/san/igroups/{igroup.uuid}/initiators`` endpoint. See [xref:{relative_path}post-protocols-san-igroups-initiators\(#-san-igroup-initiator-create\)anddelete-protocols-san-igroups-initiators-.html](#)<</SAN/igroup_initiator_create,POST
`/protocols/san/igroups/{igroup.uuid}/initiators>>` and `[DELETE
/protocols/san/igroups/{igroup.uuid}/initiators/{name}]` for more details.

```
|lun_maps
```

```
|array[link:#lun_maps[lun_maps]]
```

a|All LUN maps with which the initiator is associated.

If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.

There is an added cost to retrieving property values for ``lun_maps``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|name
|string
a|The name of the initiator group. Required in POST; optional in PATCH.
```

Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.

```
|os_type
|string
a|The host operating system of the initiator group. All initiators in the
group should be hosts of the same operating system. Required in POST;
optional in PATCH.
```

```
|protocol
|string
a|The protocols supported by the initiator group. This restricts the type
of initiators that can be added to the initiator group. Optional in POST;
if not supplied, this defaults to _mixed_.
```

The protocol of an initiator group cannot be changed after creation of the group.

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
```

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

```
//end collapsible .Definitions block
====

[[IDa4affeal68fc3beeed2858075817716f]]
= Create an initiator group

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/san/igroups`#

*Introduced In:* 9.6

Creates an initiator group.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the initiator
group.
* `name` - Name of the initiator group.
* `os_type` - Operating system of the initiator group's initiators.

== Recommended optional properties

* `initiators.name` - Name(s) of initiator group's initiators. This
property can be used to create the initiator group and populate it with
initiators in a single request.

== Default property values

If not specified in POST, the following default property values are
assigned.

* `protocol` - _mixed_ - Data protocol of the initiator group's
initiators.

== Learn more

* xref:{relative_path}protocols_san_igroups_endpoint_overview.html[DOC
/protocols/san/igroups]

== Parameters

[cols=5*,options=header]
|==
```

```
|Name
|Type
|In
|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.
```

* Default value:

```
|===
```

== Request Body

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|delete_on_unmap
|boolean
```

a|An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to `_false_` when the initiator group is created.

```
|initiators
```

```
|array[link:#igroup_initiator_no_records[igroup_initiator_no_records]]
```

a|The initiators that are members of the group. Optional in POST.

Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the ``/protocols/san/igroups/{igroup.uuid}/initiators`` endpoint. See [xref:{relative_path}post-protocols-san-igroups-initiators\(#-san-igroup-initiator-create\)](#) and [delete-protocols-san-igroups-initiators.html](#) <</SAN/igroup_initiator_create, POST `/protocols/san/igroups/{igroup.uuid}/initiators``>> and [DELETE `/protocols/san/igroups/{igroup.uuid}/initiators/{name}`] for more

details.

```
|lun_maps
|array[link:#lun_maps[lun_maps]]
a|All LUN maps with which the initiator is associated.
```

If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.

There is an added cost to retrieving property values for `lun_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|name
|string
a|The name of the initiator group. Required in POST; optional in PATCH.
```

Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.

```
|os_type
|string
a|The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH.
```

```
|protocol
|string
a|The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to _mixed_.
```

The protocol of an initiator group cannot be changed after creation of the group.

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
```

a|The unique identifier of the initiator group.

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "initiators": [
    {
      "igroup": {
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1"
    }
  ],
  "lun_maps": [
    {
      "logical_unit_number": 0,
      "lun": {
        "name": "lun1",
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "name": "igroup1",
  "os_type": "string",
  "protocol": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

====

== Response

Status: 201, Created

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|num_records
```

```
|integer
```

```
a|Number of records.
```

```
|records
```

```
|array[link:#igroup[igroup]]
```

```
a|
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "records": [
```

```
    {
```

```
      "initiators": [
```

```
        {
```

```
          "igroup": {
```

```
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
```

```
          },
```

```
          "name": "iqn.1998-01.com.corp.iscsi:name1"
```

```
        }
```

```
      ],
```

```
      "lun_maps": [
```

```
        {
```

```
          "logical_unit_number": 0,
```

```
          "lun": {
```

```
            "name": "lun1",
```

```
            "node": {
```

```
              "name": "node1",
```

```
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
            },
```

```
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
```

```
          }
```

```
        }
```

```
      ],
```

```

    "name": "igroup1",
    "os_type": "string",
    "protocol": "string",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  }
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 2621462
| The supplied SVM does not exist.

| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5373958
| An invalid initiator group name was supplied.

| 5373966
| An initiator group cannot be created in an SVM that is configured for NVMe.

| 5373969
| A supplied initiator name looks like an iSCSI IQN initiator, but the portions after the prefix are missing.

| 5373971
| A supplied initiator name looks like an iSCSI IQN initiator, but the date portion is invalid.

```

```
| 5373972
| A supplied initiator name looks like an iSCSI IQN initiator, but the
naming authority portion is invalid.

| 5373977
| A supplied initiator name looks like an iSCSI EUI initiator, but the
length is invalid.

| 5373978
| A supplied initiator name looks like an iSCSI EUI initiator, but the
format is invalid.

| 5373992
| A supplied initiator name was too long to be valid.

| 5373993
| A supplied initiator name did not match any valid format.

| 5374023
| An initiator group with the same name already exists.

| 5374038
| An invalid Fibre Channel WWPN was supplied.

| 5374039
| An invalid iSCSI initiator name was supplied.

| 5374732
| An initiator is already in another initiator group with a conflicting
operating system type.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#igroup_initiator_no_records]
[.api-collapsible-fifth-title]
igroup_initiator_no_records
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.
```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
|name
|string
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.
```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#lun]
[.api-collapsible-fifth-title]
lun
```

The LUN to which the initiator group is mapped.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the LUN.
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|The unique identifier of the LUN.
```

```
|===
```

```
[#lun_maps]
[.api-collapsible-fifth-title]
lun_maps
```

A LUN map with which the initiator group is associated.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|logical_unit_number
|integer
a|The logical unit number assigned to the LUN for initiators in the
initiator group.
```

```
|lun
|link:#lun[lun]
```


a|The LUN to which the initiator group is mapped.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#igroup]

[.api-collapsible-fifth-title]

igroup

An initiator group (igroup) is a collection of Fibre Channel (FC) world wide port names (WWPN), and/or iSCSI Qualified Names (IQNs), and/or iSCSI EUIs (Extended Unique Identifiers) that identify host initiators.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, create an initiator group containing the hosts' initiator names, then create a LUN map that associates the initiator group with the LUN.

An initiator can appear in multiple initiator groups. An initiator group can be mapped to multiple LUNs. A specific initiator can be mapped to a specific LUN only once.

All initiators in an initiator group must be from the same operating

system. The initiator group's operating system is specified when the initiator group is created.

When an initiator group is created, the ``protocol`` property is used to restrict member initiators to Fibre Channel (`_fcp_`), iSCSI (`_iscsi_`), or both (`_mixed_`).

Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the ``/protocols/san/igroups/{igroup.uuid}/initiators`` endpoint. See [xref:{relative_path}post-protocols-san-igroups-initiators\(#-san-igroup-initiator-create\)anddelete-protocols-san-igroups-initiators-.html](#)<</SAN/igroup_initiator_create,POST /protocols/san/igroups/{igroup.uuid}/initiators>> and [DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}] for more details.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|delete_on_unmap
```

```
|boolean
```

a|An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to `_false_` when the initiator group is created.

```
|initiators
```

```
|array[link:#igroup_initiator_no_records[igroup_initiator_no_records]]
```

a|The initiators that are members of the group. Optional in POST.

Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the ``/protocols/san/igroups/{igroup.uuid}/initiators`` endpoint. See [xref:{relative_path}post-protocols-san-igroups-initiators\(#-san-igroup-initiator-create\)anddelete-protocols-san-igroups-initiators-.html](#)<</SAN/igroup_initiator_create,POST /protocols/san/igroups/{igroup.uuid}/initiators>> and [DELETE /protocols/san/igroups/{igroup.uuid}/initiators/{name}] for more details.

```
|lun_maps
|array[link:#lun_maps[lun_maps]]
a|All LUN maps with which the initiator is associated.
```

If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.

There is an added cost to retrieving property values for `lun_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|name
|string
a|The name of the initiator group. Required in POST; optional in PATCH.
```

Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.

```
|os_type
|string
a|The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH.
```

```
|protocol
|string
a|The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to _mixed_.
```

The protocol of an initiator group cannot be changed after creation of the group.

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

|===

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID2388fc382bd074a9a3e2878ff12082ef]]
= Retrieve initiators of an initiator group

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/san/igroups/{igroup.uuid}/initiators`#

*Introduced In:* 9.6

Retrieves initiators of an initiator group.

== Related ONTAP commands

* `lun igroup show`

== Learn more

* xref:{relative_path}protocols_san_igroups_endpoint_overview.html[DOC
/protocols/san/igroups]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required

```

```

|Description

|igroup.uuid
|string
|path
|True
a|The unique identifier of the initiator group.


|fields
|array[string]
|query
|False
a|Specify the fields to return.


|max_records
|integer
|query
|False
a|Limit the number of records returned.


|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

```

```

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#igroup_initiator[igroup_initiator]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  }
}

```

```

    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "4ea7a442-86d1-11e0-a61c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1"
    }
  ]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 5374852
| The initiator group specified in the URI does not exist.
|====

```

[cols=3*,options=header]

```

|====
|Name
|Type
|Description

```

```

|error
|link:#error[error]

```



```

a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|next
|link:#href[href]
a|
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#records]
[.api-collapsible-fifth-title]
records
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.
```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
|name
|string
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.
```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

|===

```
[#igroup_initiator]
[.api-collapsible-fifth-title]
igroup_initiator
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.
```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
|name
|string
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.
```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
```

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

|target

```

|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDc0a067c43b210eb13e5cfe8164400da1]]
= Add initiators to an initiator group

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/san/igroups/{igroup.uuid}/initiators`#

*Introduced In:* 9.6

Adds one or more initiators to an initiator group.

== Required properties

* `name` or `records.name` - Initiator name(s) to add to the initiator
group.

== Related ONTAP commands

* `lun igroup add`

== Learn more

* xref:{relative_path}protocols_san_igroups_endpoint_overview.html[DOC
/protocols/san/igroups]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

```

```
|igroup.uuid
|string
|path
|True
a|The unique identifier of the initiator group.
```

```
|return_records
|boolean
|query
|False
a|The default is false. If set to true, the records are returned.
```

* Default value:

```
|===
```

== Request Body

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.
```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
|name
|string
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.
```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

```
|records
|array[link:#records[records]]
a|An array of initiators specified to add multiple initiators to an
initiator group in a single API call. Valid in POST only and not allowed
when the `name` property is used.
```

```
|===
```

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "igroup": {
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "iqn.1998-01.com.corp.iscsi:name1",
  "records": [
    {
      "igroup": {
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1"
    }
  ]
}
```

====

== Response

Status: 201, Created


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#igroup_initiator[igroup_initiator]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "records": [
    {
      "igroup": {
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1",
      "records": [
        {
          "igroup": {
            "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
          },
          "name": "iqn.1998-01.com.corp.iscsi:name1"
        }
      ]
    }
  ]
}
=====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 1254193

| Adding an initiator would cause the initiator to be mapped to the same LUN more than once.

| 1254324

| Adding an initiator would cause the initiator to have the same logical unit identifier for multiple LUN maps.

| 5373969

| A supplied initiator name looks like an iSCSI IQN initiator, but the portions after the prefix are missing.

| 5373971

| A supplied initiator name looks like an iSCSI IQN initiator, but the date portion is invalid.

| 5373972

| A supplied initiator name looks like an iSCSI IQN initiator, but the naming authority portion is invalid.

| 5373977

| A supplied initiator name looks like an iSCSI EUI initiator, but the length is invalid.

| 5373978

| A supplied initiator name looks like an iSCSI EUI initiator, but the format is invalid.

| 5373992

| A supplied initiator name was too long to be valid.

| 5373993

| A supplied initiator name did not match any valid format.

| 5374033

| Initiators must be supplied.

| 5374035

| A supplied initiator is already in the initiator group.

| 5374038

```

| An invalid Fibre Channel WWPN was supplied.

| 5374039
| An invalid iSCSI initiator name was supplied.

| 5374734
| An initiator is already in another initiator group with a conflicting
operating system type.

| 5374852
| The initiator group specified in the URI does not exist.

| 5374853
| You can add initiators to an initiator group using the `records`
property, or the `name` property, but you cannot use both in the same
request.

| 5374854
| Only `records` property elements should be populated with the `name`
property values.
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
[cols=3*,options=header]
|==
|Name
|Type
|Description

|uuid
|string
a|The unique identifier of the initiator group.
```

```
|==
```

```
[#records]
[.api-collapsible-fifth-title]
records
```

```
[cols=3*,options=header]
|==
|Name
|Type
|Description

|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.
```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
|name
|string
```

a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator. Valid in POST only and not allowed when the `records` property is used.

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

|===

```
[#igroup_initiator]
[.api-collapsible-fifth-title]
igroup_initiator
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.
```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
|name
|string
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.
```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

```
|records
|array[link:#records[records]]
a|An array of initiators specified to add multiple initiators to an
initiator group in a single API call. Valid in POST only and not allowed
when the `name` property is used.
```

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID8f10fa0dad98dc4a27e0064e0f3f28a9]]
= Delete an initiator from an initiator group

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/san/igroups/{igroup.uuid}/initiators/{name}`#

*Introduced In:* 9.6

Deletes an initiator from an initiator group.

== Related ONTAP commands

* `lun igroup remove`

== Learn more

* xref:{relative_path}protocols_san_igroups_endpoint_overview.html[DOC
/protocols/san/igroups]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required

```

|Description

|igroup.uuid

|string

|path

|True

a|The unique identifier of the initiator group.

|name

|string

|path

|True

a|The initiator name.

|allow_delete_while_mapped

|boolean

|query

|False

a|Allows deletion of an initiator from of a mapped initiator group.

Deleting an initiator from a mapped initiator group makes the LUNs to which the initiator group is mapped no longer available to the initiator. This might cause a disruption in the availability of data.

This parameter should be used with caution.

* Default value:

|===

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description


```
| 1254213
| The initiator group is mapped to one or more LUNs and
`allow_delete_while_mapped` has not been specified.

| 5374034
| The initiator is not a member of the group.

| 5374852
| The initiator group specified in the URI does not exist.
|===
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|error
|link:#error[error]
a|

|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions
```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID16940509adc31df837d6473e41641eda]]
= Retrieve an initiator
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/san/igroups/{igroup.uuid}/initiators/{name}`#
```

Introduced In: 9.6

Retrieves an initiator of an initiator group.

== Related ONTAP commands

* `lun igroup show`

== Learn more

* xref:{relative_path}protocols_san_igroups_endpoint_overview.html[DOC
/protocols/san/igroups]

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
```

```

|Required
|Description

|igroup.uuid
|string
|path
|True
a|The unique identifier of the initiator group.

|name
|string
|path
|True
a|Initiator name

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in
other initiator groups.

```

```
|name
|string
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.
```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

```
|===
```

.Example response

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroup": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "iqn.1998-01.com.corp.iscsi:name1"
}
```

```
=====
```

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 4
| The initiator is not a member of the initiator group.

| 5374852
| The initiator group specified in the URI does not exist.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

```
.Example error
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```
=====
```

```
== Definitions
```

```
[.api-def-first-level]
```

```
.See Definitions
```

```
[%collapsible%closed]
```

```
//Start collapsible Definitions block
```

```
====
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
[cols=3*,options=header]
|===
```

```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The unique identifier of the initiator group.

```

```

|===

```

```

[#records]
[.api-collapsible-fifth-title]
records

```

```

[cols=3*,options=header]

```

```

|===

```

```

|Name
|Type
|Description

```

```

|_links
|link:#_links[_links]
a|

```

```

|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.

```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```

|name
|string
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.

```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._` prefix followed by 16 hexadecimal characters.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]

|===

|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]

|===

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

```

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDc7bd2ce9f4d0fb141f58414bcdd19a20]]
= Delete an initiator group

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/san/igroups/{uuid}`#

*Introduced In:* 9.6

Deletes an initiator group.

== Related ONTAP commands

* `lun igroup delete`

== Learn more

* xref:{relative_path}protocols_san_igroups_endpoint_overview.html[DOC
/proTOCOLS/san/igroups]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In

```

```
|Required
|Description

|uuid
|string
|path
|True
a|The unique identifier of the initiator group.
```

```
|allow_delete_while_mapped
|boolean
|query
|False
a|Allows deletion of a mapped initiator group.
```

Deleting a mapped initiator group makes the LUNs to which the initiator group is mapped no longer available. This might cause a disruption in the availability of data.

This parameter should be used with caution.

* Default value:

```
|===
```

```
== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error Code | Description
```

```
| 1254213
```

```
| The initiator group is mapped to one or more LUNs and
`allow_delete_while_mapped` has not been specified.
```

```
| 5374852
```

```
| The initiator group does not exist.
```

```

|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

```

|===

//end collapsible .Definitions block
=====

[[ID98eaf11841f94225de82ddf887ccf65e]]
= Retrieve an initiator group

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/san/igroups/{uuid}`#

Introduced In: 9.6

Retrieves an initiator group.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `lun_maps.+++`

== Related ONTAP commands

* `lun igroup show`
* `lun mapping show`

== Learn more

* [xref:{relative_path}protocols_san_igroups_endpoint_overview.html\[DOC /protocols/san/igroups\]](#)

== Parameters

[cols=5*,options=header]
|===

|Name

```

|Type
|In
|Required
|Description

|uuid
|string
|path
|True
a|The unique identifier of the initiator group.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|delete_on_unmap
|boolean
a|An option that causes the initiator group to be deleted when the last
LUN map associated with it is deleted. Optional in POST and PATCH. This
property defaults to _false_ when the initiator group is created.

|initiators
|array[link:#igroup_initiator_no_records[igroup_initiator_no_records]]
a|The initiators that are members of the group. Optional in POST.

Zero or more initiators can be supplied when the initiator group is

```

created. After creation, initiators can be added or removed from the initiator group using the `/protocols/san/igroups/{igroup.uuid}/initiators` endpoint. See xref:{relative_path}post-protocols-san-igroups-initiators\(#-san-igroup-initiator-create\) and delete-protocols-san-igroups-initiators-.html` `<</SAN/igroup_initiator_create,POST` `/protocols/san/igroups/{igroup.uuid}/initiators>>` and `[DELETE` `/protocols/san/igroups/{igroup.uuid}/initiators/{name}]` for more details.

`|lun_maps`
`|array[link:#lun_maps[lun_maps]]`
`a|All LUN maps with which the initiator is associated.`

If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.

There is an added cost to retrieving property values for ``lun_maps``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

`|name`
`|string`
`a|The name of the initiator group. Required in POST; optional in PATCH.`

Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.

`|os_type`
`|string`
`a|The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH.`

`|protocol`
`|string`
`a|The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group. Optional in POST; if not supplied, this defaults to _mixed_.`

The protocol of an initiator group cannot be changed after creation of the group.


```
|svm
|link:#svm[svm]
a|

|uuid
|string
a|The unique identifier of the initiator group.

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "initiators": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "name": "iqn.1998-01.com.corp.iscsi:name1"
    }
  ],
  "lun_maps": [
    {
      "_links": {
        "self": {
```

```

        "href": "/api/resourcelink"
      }
    },
    "logical_unit_number": 0,
    "lun": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "lun1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "name": "igroup1",
  "os_type": "string",
  "protocol": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 5374852
| The initiator group does not exist.
|===

```

```

[cols=3*,options=header]

```

```

|===

```

```

|Name

```

```

|Type

```

```

|Description

```

```

|error

```

```

|link:#error[error]

```

```

a|

```

```

|===

```

```

.Example error

```

```

[%collapsible%closed]

```

```

=====

```

```

[source,json,subs=+macros]

```

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

```

=====

```

```

== Definitions

```

```

[.api-def-first-level]

```

```

.See Definitions

```

```

[%collapsible%closed]

```

```
//Start collapsible Definitions block
```

```
====
```

```
[#href]
```

```
[.api-collapsible-fifth-title]
```

```
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|href
```

```
|string
```

```
a|
```

```
|===
```

```
[#_links]
```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:href[href]
```

```
a|
```

```
|===
```

```
[#igroup]
```

```
[.api-collapsible-fifth-title]
```

```
igroup
```

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The unique identifier of the initiator group.

```

```

|===

[#igroup_initiator_no_records]
[.api-collapsible-fifth-title]
igroup_initiator_no_records

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.

```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```

|name
|string
a|The FC WWPN, iSCSI IQN, or iSCSI EUI that identifies the host initiator.
Valid in POST only and not allowed when the `records` property is used.

```

An FC WWPN consist of 16 hexadecimal digits grouped as 8 pairs separated by colons. The format for an iSCSI IQN is `_iqn.yyyy-mm.reverse_domain_name:any_`. The iSCSI EUI format consists of the `_eui._`

prefix followed by 16 hexadecimal characters.

|===

[#node]
[.api-collapsible-fifth-title]
node

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#lun]
[.api-collapsible-fifth-title]
lun

The LUN to which the initiator group is mapped.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```
|name
|string
a|The name of the LUN.
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|The unique identifier of the LUN.
```

```
|===
```

```
[#lun_maps]
[.api-collapsible-fifth-title]
lun_maps
```

A LUN map with which the initiator group is associated.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
 |_links
|link:#_links[_links]
a|
```

```
|logical_unit_number
|integer
a|The logical unit number assigned to the LUN for initiators in the
initiator group.
```

```
|lun
|link:#lun[lun]
a|The LUN to which the initiator group is mapped.
```

```
|===
```

```

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```



```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDdec60086ccc802b748d530a71a6cdf82]]
= Update an initiator group

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
```

block]#`/protocols/san/igroups/{uuid}`#

Introduced In: 9.6

Updates an initiator group.

== Related ONTAP commands

* `lun igroup modify`

* `lun igroup rename`

== Learn more

* xref:{relative_path}protocols_san_igroups_endpoint_overview.html[DOC
/protocols/san/igroups]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|uuid
|string
|path
|True

a|The unique identifier of the initiator group.

|===

== Request Body

[cols=3*,options=header]
|===

|Name
|Type
|Description

|delete_on_unmap

|boolean

a|An option that causes the initiator group to be deleted when the last LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to `_false_` when the initiator group is created.

|lun_maps

|array[link:#lun_maps[lun_maps]]

a|All LUN maps with which the initiator is associated.

If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.

There is an added cost to retrieving property values for ``lun_maps``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|name

|string

a|The name of the initiator group. Required in POST; optional in PATCH.

Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.

|os_type

|string

a|The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH.

|svm

|link:#svm[svm]

a|

|uuid

|string

a|The unique identifier of the initiator group.

|===

.Example request

```
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "lun_maps": [
    {
      "logical_unit_number": 0,
      "lun": {
        "name": "lun1",
        "node": {
          "name": "node1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "name": "igroup1",
  "os_type": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
====

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|====
```

```
| Error Code | Description
```

```
| 5373958
```

```
| An invalid initiator group name was supplied for a rename operation.
```

```
| 5374023
```

```
| A rename operation failed because an initiator group with the same name
```

already exists.

| 5374733
| An initiator is already in another initiator group with a conflicting operating system type.

| 5374852
| The initiator group does not exist.

| 5374868
| The initiator group was partially modified before an error was encountered while renaming the initiator group.
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#igroup]
[.api-collapsible-fifth-title]
igroup

The initiator group in which the initiator is found.

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#igroup_initiator_no_records]
[.api-collapsible-fifth-title]
igroup_initiator_no_records
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|igroup
|link:#igroup[igroup]
a|The initiator group in which the initiator is found.
```

Note that this does not mean that the initiator cannot also be found in other initiator groups.

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

```
[cols=3*,options=header]
|===
|Name
```

```
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#lun]
[.api-collapsible-fifth-title]
lun
```

The LUN to which the initiator group is mapped.

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the LUN.
```

```
|node
|link:#node[node]
a|
```

```
|uuid
|string
a|The unique identifier of the LUN.
```

```
|===
```

```
[#lun_maps]
[.api-collapsible-fifth-title]
lun_maps
```

A LUN map with which the initiator group is associated.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|logical_unit_number
```

```
|integer
```

a|The logical unit number assigned to the LUN for initiators in the initiator group.

```
|lun
```

```
|link:#lun[lun]
```

a|The LUN to which the initiator group is mapped.

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

svm

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|The name of the SVM.

```
|uuid
```

```
|string
```

a|The unique identifier of the SVM.

```
|===
```



```
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

An initiator group (igroup) is a collection of Fibre Channel (FC) world wide port names (WWPN), and/or iSCSI Qualified Names (IQNs), and/or iSCSI EUIs (Extended Unique Identifiers) that identify host initiators.

Initiator groups are used to control which hosts can access specific LUNs. To grant access to a LUN from one or more hosts, create an initiator group containing the hosts' initiator names, then create a LUN map that associates the initiator group with the LUN.

An initiator can appear in multiple initiator groups. An initiator group can be mapped to multiple LUNs. A specific initiator can be mapped to a specific LUN only once.

All initiators in an initiator group must be from the same operating system. The initiator group's operating system is specified when the initiator group is created.

When an initiator group is created, the `protocol` property is used to restrict member initiators to Fibre Channel (`_fcp_`), iSCSI (`_iscsi_`), or both (`_mixed_`).

Zero or more initiators can be supplied when the initiator group is created. After creation, initiators can be added or removed from the initiator group using the ``/protocols/san/igroups/{igroup.uuid}/initiators`` endpoint. See `xref:{relative_path}post-protocols-san-igroups-initiators(#-san-igroup-initiator-create)anddelete-protocols-san-igroups-initiators-.html<</SAN/igroup_initiator_create,POST` `/protocols/san/igroups/{igroup.uuid}/initiators>>` and `[DELETE` `/protocols/san/igroups/{igroup.uuid}/initiators/{name}]` for more details.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|delete_on_unmap
```

```
|boolean
```

a|An option that causes the initiator group to be deleted when the last

LUN map associated with it is deleted. Optional in POST and PATCH. This property defaults to `_false_` when the initiator group is created.

```
|lun_maps
|array[link:#lun_maps[lun_maps]]
a|All LUN maps with which the initiator is associated.
```

If the requested igroup is part of a remote, non-local, MetroCluster SVM, the LUN maps are not retrieved.

There is an added cost to retrieving property values for ``lun_maps``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|name
|string
a|The name of the initiator group. Required in POST; optional in PATCH.
```

Note that renaming an initiator group must be done in a PATCH request separate from any other modifications.

```
|os_type
|string
a|The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system. Required in POST; optional in PATCH.
```

```
|svm
|link:#svm[svm]
a|
```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

:leveloffset: -1

= Manage iSCSI credentials

:leveloffset: +1

[[ID0ba261ad2d4612c686ef6ea765d6334f]]

= Protocols SAN iSCSI credentials endpoint overview

== Overview

An iSCSI credentials object defines authentication credentials to be used between an initiator and ONTAP. It identifies an authentication type, user names, and passwords that must be used to authenticate a specific initiator.

The iSCSI credentials REST API allows you to create, update, delete, and discover iSCSI credentials.

== How iSCSI authentication works

An iSCSI credentials object defines the authentication credentials to be used between an initiator and ONTAP. While establishing an iSCSI connection, the initiator sends a login request to ONTAP to begin an iSCSI session. ONTAP then either permits or denies the login request, or determines that a login is not required.

For an initiator, you can specify an authentication type, user names and passwords, and a whitelist of optional network addresses from which the initiator is allowed to connect.

== iSCSI authentication methods

* Challenge-Handshake Authentication Protocol (CHAP) - The initiator logs in using a CHAP user name and password. There are two types of CHAP user names and passwords:

** Inbound - ONTAP authenticates the initiator. Inbound settings are required if you are using CHAP authentication.

** Outbound - These are optional credentials to enable the initiator to authenticate ONTAP. You can use credentials only if inbound credentials are also being used.

* deny - The initiator is denied access to ONTAP.

* none - ONTAP does not require authentication for the initiator.

The CHAP inbound/outbound password can be any valid string or an even number of valid hexadecimal digits preceded by '0X' or '0x'.

== Initiator address list

The initiator address list is a way to specify valid IP addresses from which the initiator is allowed to connect. If the list is specified and the source address of an iSCSI connection is not in the list, the connection is rejected. Initiator addresses can be specified in either IPv4 or IPv6 format and in one of two forms:

* Range

```
----  
{  
  "start": "192.168.0.0",  
  "end": "192.168.0.255"  
}  
----
```

* Mask

```
----  
{  
  "address": "192.168.0.0",  
  "netmask": "24"  
}  
----
```

== Initiator "default"

The default iSCSI authentication definition is created when the iSCSI service is created. An iSCSI credentials object with `_default_` as the initiator name identifies the default authentication for an SVM. The default credentials are used for any initiator that does not have specific

iSCSI credentials. The default iSCSI authentication method is `_none_`, but can be changed to `_deny_` or `_CHAP_`. The default credentials object does not support an initiator address list.

== Examples

=== Creating iSCSI credentials requiring no authentication

The API:

POST /api/protocols/san/iscsi/credentials

The call:

```
curl -X POST 'https://<mgmt-ip>/api/protocols/san/iscsi/credentials' -H
'accept: application/hal+json' -d '{ "svm": { "name": "svm1" },
"initiator": "iqn.1992-08.com.netapp:initiator1", "authentication_type":
"none" }'
```

'''

=== Creating iSCSI credentials using CHAP inbound authentication

The API:

POST /api/protocols/san/iscsi/credentials

The call:

```
curl -X POST 'https://<mgmt-ip>/api/protocols/san/iscsi/credentials' -H
'accept: application/hal+json' -d '{ "svm": { "name": "svm1" },
"initiator": "iqn.1992-08.com.netapp:initiator2", "authentication_type":
"CHAP", "chap": { "inbound": { "user": "user1", "password": "password1" }
} }'
```

'''

=== Retrieving all properties of all iSCSI credentials

The ``fields`` query parameter is used to request all iSCSI credentials properties.

Passwords are not included in the GET output.

```

# The API:
GET /api/protocols/san/iscsi/credentials

# The call:
curl -X GET 'https://<mgmt-
ip>/api/protocols/san/iscsi/credentials?fields=*' -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
          }
        }
      },
      "initiator": "default",
      "authentication_type": "none",
      "_links": {
        "self": {
          "href": "/api/protocols/san/iscsi/credentials/19d04b8e-94d7-11e8-
8370-005056b48fd2/default"
        }
      }
    },
    {
      "svm": {
        "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
          }
        }
      },
      "initiator": "iqn.1992-08.com.netapp:initiator1",
      "authentication_type": "none",
      "_links": {
        "self": {
          "href": "/api/protocols/san/iscsi/credentials/19d04b8e-94d7-11e8-

```

```

8370-005056b48fd2/iqn.1992-08.com.netapp:initiator1"
    }
  },
  {
    "svm": {
      "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "initiator": "iqn.1992-08.com.netapp:initiator2",
    "authentication_type": "chap",
    "chap": {
      "inbound": {
        "user": "user1"
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/credentials/19d04b8e-94d7-11e8-8370-005056b48fd2/iqn.1992-08.com.netapp:initiator2"
      }
    }
  },
  {
    "svm": {
      "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "initiator": "default",
    "authentication_type": "none",
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-005056b48fd2/default"
      }
    }
  }
}

```



```

    },
    {
      "svm": {
        "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
        "name": "svm2",
        "_links": {
          "self": {
            "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
          }
        }
      },
      "initiator": "iqn.1992-08.com.netapp:initiator2",
      "authentication_type": "none",
      "_links": {
        "self": {
          "href": "/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-005056b48fd2/iqn.1992-08.com.netapp:initiator2"
        }
      }
    },
    {
      "svm": {
        "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
        "name": "svm2",
        "_links": {
          "self": {
            "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
          }
        }
      },
      "initiator": "iqn.1992-08.com.netapp:initiator3",
      "authentication_type": "deny",
      "_links": {
        "self": {
          "href": "/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-005056b48fd2/iqn.1992-08.com.netapp:initiator3"
        }
      }
    }
  ],
  "num_records": 6,
  "_links": {
    "self": {
      "href": "/api/protocols/san/iscsi/credentials?fields=*"
    }
  }
}

```

```

}
----

'''

=== Retrieving specific iSCSI credentials

----

# The API:
GET /api/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}

# The call:
curl -X GET 'https://<mgmt-
ip>/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-
005056b48fd2/iqn.1992-08.com.netapp:initiator2' -H 'accept:
application/hal+json'

# The response:
{
  "svm": {
    "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
    "name": "svm2",
    "_links": {
      "self": {
        "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
      }
    }
  },
  "initiator": "iqn.1992-08.com.netapp:initiator2",
  "authentication_type": "chap",
  "chap": {
    "inbound": {
      "user": "user1"
    }
  },
  "_links": {
    "self": {
      "href": "/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-
005056b48fd2/iqn.1992-08.com.netapp:initiator2"
    }
  }
}
----

'''

```

```
=== Updating the authentication type of iSCSI credentials
```

```
----
```

```
# The API:
```

```
PATCH /api/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}
```

```
# The call:
```

```
curl -X PATCH 'https://<mgmt-  
ip>/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-  
005056b48fd2/iqn.1992-08.com.netapp:initiator2' -H 'accept:  
application/hal+json' -d '{ "authentication_type": "chap", "chap": {  
"inbound": { "user": "user1", "password": "password1" } } }'
```

```
----
```

```
'''
```

```
=== Updating the initiator address list of iSCSI credentials
```

```
----
```

```
# The API:
```

```
PATCH /api/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}
```

```
# The call:
```

```
curl -X PATCH 'https://<mgmt-  
ip>/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-  
005056b48fd2/iqn.1992-08.com.netapp:initiator2' -H 'accept:  
application/hal+json' -d '{ "initiator_address": { "ranges": [ { "start":  
"192.168.0.0", "end": "192.168.255.255" } ] } }'
```

```
----
```

```
'''
```

```
=== Deleting iSCSI credentials
```

```
----
```

```
# The API:
```

```
DELETE /api/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}
```

```
# The call:
```

```
curl -X DELETE 'https://<mgmt-  
ip>/api/protocols/san/iscsi/credentials/25f617cf-94d7-11e8-8370-  
005056b48fd2/iqn.1992-08.com.netapp:initiator2' -H 'accept:  
application/hal+json'
```

```
----
```

[[ID6117167701f59e2a3dd12a05a31b9dba]]

= Retrieve iSCSI credentials

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/san/iscsi/credentials`#

Introduced In: 9.6

Retrieves iSCSI credentials.

== Related ONTAP commands

* `vserver iscsi security show`

== Learn more

*

xref:{relative_path}protocols_san_iscsi_credentials_endpoint_overview.html
[DOC /protocols/san/iscsi/credentials]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|initiator

|string

|query

|False

a|Filter by initiator

|svm.uuid

|string

|query

```
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|chap.inbound.user
|string
|query
|False
a|Filter by chap.inbound.user

|chap.outbound.user
|string
|query
|False
a|Filter by chap.outbound.user

|initiator_address.masks.address
|string
|query
|False
a|Filter by initiator_address.masks.address

|initiator_address.masks.family
|string
|query
|False
a|Filter by initiator_address.masks.family

|initiator_address.masks.netmask
|string
|query
|False
a|Filter by initiator_address.masks.netmask

|initiator_address.ranges.family
```

```
|string
|query
|False
a|Filter by initiator_address.ranges.family
```

```
|initiator_address.ranges.end
|string
|query
|False
a|Filter by initiator_address.ranges.end
```

```
|initiator_address.ranges.start
|string
|query
|False
a|Filter by initiator_address.ranges.start
```

```
|authentication_type
|string
|query
|False
a|Filter by authentication_type
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

|return_timeout
|integer
|query
|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by
|array[string]
|query
|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name
|Type
|Description

|_links
|link:#_links[_links]

a|

|num_records
|integer
a|Number of records.

```
|records
|array[link:#iscsi_credentials[iscsi_credentials]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_type": "string",
      "chap": {
        "inbound": {
          "user": "string"
        },
        "outbound": {
          "user": "string"
        }
      },
      "initiator": "iqn.1998-01.com.corp.iscsi:name1",
      "initiator_address": {
        "masks": [
          {
            "address": "10.10.10.7",
            "family": "string",
            "netmask": "24"
          }
        ],
        "ranges": [
```



```

        {
            "end": "10.10.10.7",
            "family": "string",
            "start": "10.10.10.7"
        }
    ]
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
    "error": {
        "arguments": [

```

```

    {
      "code": "string",
      "message": "string"
    }
  ],
  "code": "4",
  "message": "entry doesn't exist",
  "target": "uuid"
}

```

====

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

====

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

|===

```

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===

[#inbound]
[.api-collapsible-fifth-title]
inbound

Inbound CHAP credentials.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|user
|string
a|The inbound CHAP user name. Optional in POST and PATCH.

|===

```

```
[#outbound]
[.api-collapsible-fifth-title]
outbound
```

Output CHAP credentials.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|user
|string
a|The outbound CHAP user name. Optional in POST and PATCH.

|===
```

```
[#chap]
[.api-collapsible-fifth-title]
chap
```

Challenge-Handshake Authentication Protocol (CHAP) credentials.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|inbound
|link:#inbound[inbound]
a|Inbound CHAP credentials.

|outbound
|link:#outbound[outbound]
a|Output CHAP credentials.

|===
```

```
[#ip_info]
[.api-collapsible-fifth-title]
ip_info
```

IP information

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|address
```

```
|string
```

```
a|IPv4 or IPv6 address
```

```
|family
```

```
|string
```

```
a|IPv4 or IPv6
```

```
|netmask
```

```
|string
```

```
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.
```

```
|===
```

```
[#ip_address_range]
```

```
[.api-collapsible-fifth-title]
```

```
ip_address_range
```

IP address range

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|end
|string
a|IPv4 or IPv6 address
```

```
|family
|string
a|IPv4 or IPv6
```

```
|start
|string
a|IPv4 or IPv6 address
```

```
|===
```

```
[#initiator_address]
[.api-collapsible-fifth-title]
initiator_address
```

Initiator address ranges.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|masks
|array[link:#ip_info[ip_info]]
a|
```

```
|ranges
|array[link:#ip_address_range[ip_address_range]]
a|
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#iscsi_credentials]
[.api-collapsible-fifth-title]
iscsi_credentials

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_type
|string
a|The iSCSI authentication type. Required in POST and optional in PATCH.

|chap
|link:#chap[chap]
a|Challenge-Handshake Authentication Protocol (CHAP) credentials.

```

```

|initiator
|string
a|The iSCSI initiator to which the credentials apply. Required in POST.

|initiator_address
|link:#initiator_address[initiator_address]
a|Initiator address ranges.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name

```



```

|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

[[IDf0d66635311f8d32c56f2769c67bab2a]]
= Create iSCSI credentials

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/san/iscsi/credentials`#

*Introduced In:* 9.6

Creates iSCSI credentials.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the iSCSI
credentials.
* `initiator` - Initiator for which the iSCSI credentials are to be
created.

```

* `authentication_type` - Type of authentication to use for the credentials.

== Recommended optional properties

* `chap.inbound.user` - In-bound CHAP authentication user name.

* `chap.inbound.password` - In-bound CHAP authentication password.

* `chap.outbound.user` - Out-bound CHAP authentication user name.

* `chap.outbound.password` - Out-bound CHAP authentication password.

== Related ONTAP commands

* `vserver iscsi security create`

== Learn more

*

xref:{relative_path}protocols_san_iscsi_credentials_endpoint_overview.html
[DOC /protocols/san/iscsi/credentials]

== Parameters

[cols=5*,options=header]
|==

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|==

== Request Body

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

|authentication_type
|string
a|The iSCSI authentication type. Required in POST and optional in PATCH.

|chap
|link:#chap[chap]
a|Challenge-Handshake Authentication Protocol (CHAP) credentials.

|initiator
|string
a|The iSCSI initiator to which the credentials apply. Required in POST.

|initiator_address
|link:#initiator_address[initiator_address]
a|Initiator address ranges.

|svm
|link:#svm[svm]
a|

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "authentication_type": "string",
  "chap": {
    "inbound": {
      "password": "string",
      "user": "string"
    },
    "outbound": {
      "password": "string",
      "user": "string"
    }
  }
}

```

```

},
"initiator": "iqn.1998-01.com.corp.iscsi:name1",
"initiator_address": {
  "masks": [
    {
      "address": "10.10.10.7",
      "netmask": "24"
    }
  ],
  "ranges": [
    {
      "end": "10.10.10.7",
      "start": "10.10.10.7"
    }
  ]
},
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
=====

```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#iscsi_credentials[iscsi_credentials]]
a|

|===

```

.Example response

```

[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "authentication_type": "string",
      "chap": {
        "inbound": {
          "password": "string",
          "user": "string"
        },
        "outbound": {
          "password": "string",
          "user": "string"
        }
      },
      "initiator": "iqn.1998-01.com.corp.iscsi:name1",
      "initiator_address": {
        "masks": [
          {
            "address": "10.10.10.7",
            "netmask": "24"
          }
        ],
        "ranges": [
          {
            "end": "10.10.10.7",
            "start": "10.10.10.7"
          }
        ]
      },
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 2621462

| An SVM with the specified UUID does not exist.

| 2621706

| Both the SVM UUID and SVM name were supplied, but they do not refer to the same SVM.

| 2621707

| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5373969

| A non-empty qualifier is required after the prefix. An example of a valid IQN is `_iqn.1995-08.com.example:string_`.

| 5373970

| The IQN prefix is invalid. The correct IQN prefix is `_iqn_`. An example of a valid IQN is `_iqn.1995-08.com.example:string_`.

| 5373971

| The date field is invalid. A valid date field is `_yyyy-mm_`. An example of a valid IQN is `_iqn.1995-08.com.example:string_`.

| 5373972

| The naming authority and string fields can contain only the characters `_a-z_`, `_0-9_`, `_.`, `_:`, and `_-`.

| 5373977

| The EUI-64 identifier field must be exactly 16 hexadecimal digits.

| 5373978

| The EUI formatted initiator name supplied is invalid. A valid EUI format is `_eui.XXXXXXXXXXXXXXXXXX_`, where X is a hexadecimal digit.

| 5373997

| The initiator name supplied is invalid. The valid initiator name formats are `_iqn.1995-08.com.example:string_` or `_eui.0123456789abcdef_`.

| 5374078

| The iSCSI service does not exist.

| 5374142

| An iSCSI security credential already exists for the specified initiator.

| 5374145
 | The iSCSI security password must contain an even number of valid hex digits.

| 5374147
 | The CHAP inbound and outbound passwords must be different.

| 5374149
 | The inbound user and password properties are required for CHAP authentication.

| 5374150
 | Outbound CHAP authentication requires an outbound password.

| 5374855
 | The value for property `initiator_address.ranges.start` is greater than the value for property `initiator_address.ranges.end`.

| 5374856
 | The value for property `initiator_address.ranges.start` does not belong to the same IP address family as the value for property `initiator_address.ranges.end`.

| 5374900
 | Setting the CHAP authentication properties are not supported with authentication types `_none_` or `_deny_`.

|===

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#inbound]
[.api-collapsible-fifth-title]
inbound

Inbound CHAP credentials.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|password
|string
a|The inbound CHAP password. Write-only; optional in POST and PATCH.

|user
|string
a|The inbound CHAP user name. Optional in POST and PATCH.

|===

[#outbound]
[.api-collapsible-fifth-title]
outbound

Output CHAP credentials.

[cols=3*,options=header]
|===
|Name
|Type

```



```

|Description

|password
|string
a|The outbound CHAP password. Write-only; optional in POST and PATCH.

|user
|string
a|The outbound CHAP user name. Optional in POST and PATCH.

|===

[#chap]
[.api-collapsible-fifth-title]
chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|inbound
|link:#inbound[inbound]
a|Inbound CHAP credentials.

|outbound
|link:#outbound[outbound]
a|Output CHAP credentials.

|===

[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.

|===

[#ip_address_range]
[.api-collapsible-fifth-title]
ip_address_range

IP address range

[cols=3*,options=header]
|===
|Name
|Type
|Description

|end
|string
a|IPv4 or IPv6 address

|start
|string
a|IPv4 or IPv6 address

|===

```

```
[#initiator_address]
[.api-collapsible-fifth-title]
initiator_address
```

Initiator address ranges.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|masks
|array[link:#ip_info[ip_info]]
a|

|ranges
|array[link:#ip_address_range[ip_address_range]]
a|

|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===
```

```

[#iscsi_credentials]
[.api-collapsible-fifth-title]
iscsi_credentials

[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication_type
|string
a|The iSCSI authentication type. Required in POST and optional in PATCH.

|chap
|link:#chap[chap]
a|Challenge-Handshake Authentication Protocol (CHAP) credentials.

|initiator
|string
a|The iSCSI initiator to which the credentials apply. Required in POST.

|initiator_address
|link:#initiator_address[initiator_address]
a|Initiator address ranges.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====

[[ID90ec7299857de82499465ac618cc9044]]
= Delete iSCSI credentials

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}`#

*Introduced In:* 9.6

Deletes specified iSCSI credentials.

== Related ONTAP commands

* `vserver iscsi security delete`

== Learn more

*
xref:{relative_path}protocols_san_iscsi_credentials_endpoint_overview.html
[DOC /protocols/san/iscsi/credentials]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|The unique identifier of an SVM.

|initiator
|string
|path
```

```
|True
a|The iSCSI initiator of the credentials object.

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description

| 2621462
| An SVM with the specified UUID does not exist.

| 2621706
| Both the SVM UUID and SVM name were supplied, but they do not refer to
the same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5374148
| The default security credential cannot be deleted for an SVM.

| 5374895
| The iSCSI security credential does not exist on the specified SVM.
|===
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```

a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```



```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID51206e2ff652ed67c46c4c569319813d]]

```

= Retrieve specific iSCSI credentials

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}`#

Introduced In: 9.6

Retrieves specified iSCSI credentials.

== Related ONTAP commands

* `vserver iscsi security show`

== Learn more

*

xref:{relative_path}protocols_san_iscsi_credentials_endpoint_overview.html
[DOC /protocols/san/iscsi/credentials]

== Parameters

[cols=5*,options=header]
|==

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|The unique identifier of an SVM.

|initiator
|string
|path
|True
a|The iSCSI initiator of the credentials object.

|fields
|array[string]

```
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|authentication_type
|string
a|The iSCSI authentication type. Required in POST and optional in PATCH.

|chap
|link:#chap[chap]
a|Challenge-Handshake Authentication Protocol (CHAP) credentials.

|initiator
|string
a|The iSCSI initiator to which the credentials apply. Required in POST.

|initiator_address
|link:#initiator_address[initiator_address]
a|Initiator address ranges.

|svm
|link:#svm[svm]
a|

|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_type": "string",
  "chap": {
    "inbound": {
      "user": "string"
    },
    "outbound": {
      "user": "string"
    }
  },
  "initiator": "iqn.1998-01.com.corp.iscsi:name1",
  "initiator_address": {
    "masks": [
      {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24"
      }
    ],
    "ranges": [
      {
        "end": "10.10.10.7",
        "family": "string",
        "start": "10.10.10.7"
      }
    ]
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

```

}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#inbound]
[.api-collapsible-fifth-title]
inbound

Inbound CHAP credentials.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|user
|string
a|The inbound CHAP user name. Optional in POST and PATCH.
```

```
|===
```

```
[#outbound]
[.api-collapsible-fifth-title]
outbound
```

Output CHAP credentials.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|user
|string
a|The outbound CHAP user name. Optional in POST and PATCH.
```

```
|===
```

```
[#chap]
[.api-collapsible-fifth-title]
chap
```

Challenge-Handshake Authentication Protocol (CHAP) credentials.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|inbound
|link:#inbound[inbound]
a|Inbound CHAP credentials.
```

```
|outbound
|link:#outbound[outbound]
a|Output CHAP credentials.
```

```
|===
```

```
[#ip_info]
[.api-collapsible-fifth-title]
ip_info
```

IP information

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|address
|string
a|IPv4 or IPv6 address
```

```
|family
|string
a|IPv4 or IPv6
```

```
|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.
```

```
|===
```

```
[#ip_address_range]
[.api-collapsible-fifth-title]
ip_address_range
```

IP address range


```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|end
```

```
|string
```

```
a|IPv4 or IPv6 address
```

```
|family
```

```
|string
```

```
a|IPv4 or IPv6
```

```
|start
```

```
|string
```

```
a|IPv4 or IPv6 address
```

```
|===
```

```
[#initiator_address]
```

```
[.api-collapsible-fifth-title]
```

```
initiator_address
```

Initiator address ranges.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|masks
```

```
|array[link:#ip_info[ip_info]]
```

```
a|
```

```
|ranges
```

```
|array[link:#ip_address_range[ip_address_range]]
```

```
a|
```

```
|===
```

```

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDa2b111891902654537fb8d3869e1bf8d]]
= Update iSCSI credentials

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
```

```

block]#`/protocols/san/iscsi/credentials/{svm.uuid}/{initiator}`#

*Introduced In:* 9.6

Updates specified iSCSI credentials.

== Related ONTAP commands

* `vserver iscsi security add-initiator-address-ranges`
* `vserver iscsi security default`
* `vserver iscsi security modify`
* `vserver iscsi security remove-initiator-address-ranges`

== Learn more

*
xref:{relative_path}protocols_san_iscsi_credentials_endpoint_overview.html
[DOC /protocols/san/iscsi/credentials]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|The unique identifier of an SVM.

|initiator
|string
|path
|True
a|The iSCSI initiator of the credentials object.

|add_initiator_addresses
|boolean

```

```

|query
|False
a|If _true_, the initiator addresses in the body merge into the existing
addresses in the iSCSI security object rather than replace the existing
addresses.

* Default value:

|remove_initiator_addresses
|boolean
|query
|False
a|If _true_, the initiator addresses in the body are removed from the
existing addresses in the iSCSI security object rather than replace the
existing addresses.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|authentication_type
|string
a|The iSCSI authentication type. Required in POST and optional in PATCH.

|chap
|link:#chap[chap]
a|Challenge-Handshake Authentication Protocol (CHAP) credentials.

|initiator_address
|link:#initiator_address[initiator_address]
a|Initiator address ranges.

|svm

```

```
|link:#svm[svm]
a|

|===

.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "authentication_type": "string",
  "chap": {
    "inbound": {
      "password": "string",
      "user": "string"
    },
    "outbound": {
      "password": "string",
      "user": "string"
    }
  },
  "initiator_address": {
    "masks": [
      {
        "address": "10.10.10.7",
        "netmask": "24"
      }
    ],
    "ranges": [
      {
        "end": "10.10.10.7",
        "start": "10.10.10.7"
      }
    ]
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

== Response
```

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 2621462

| An SVM with the specified UUID does not exist.

| 2621706

| Both the SVM UUID and SVM name were supplied, but they do not refer to the same SVM.

| 2621707

| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5374145

| The iSCSI security password must contain an even number of valid hex digits.

| 5374147

| The CHAP inbound and outbound passwords must be different.

| 5374149

| The inbound user and password properties are required for CHAP authentication.

| 5374150

| Outbound CHAP authentication requires an outbound password.

| 5374155

| The functionality is not supported for the default security credential.

| 5374855

| The value for property `initiator_address.ranges.start` is greater than the value for property `initiator_address.ranges.end`.

| 5374856

| The value for property `initiator_address.ranges.start` does not belong to the same IP address family as the value for property `initiator_address.ranges.end`.

```
| 5374895
| The iSCSI security credential does not exist on the specified SVM.

| 5374900
| Setting the CHAP authentication properties are not supported with
authentication types _none_ or _deny_.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#inbound]
[.api-collapsible-fifth-title]
inbound
```

Inbound CHAP credentials.

```
[cols=3*,options=header]
|===
|Name
|Type
```



```

|Description

|password
|string
a|The inbound CHAP password. Write-only; optional in POST and PATCH.

|user
|string
a|The inbound CHAP user name. Optional in POST and PATCH.

|===

[#outbound]
[.api-collapsible-fifth-title]
outbound

Output CHAP credentials.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|password
|string
a|The outbound CHAP password. Write-only; optional in POST and PATCH.

|user
|string
a|The outbound CHAP user name. Optional in POST and PATCH.

|===

[#chap]
[.api-collapsible-fifth-title]
chap

Challenge-Handshake Authentication Protocol (CHAP) credentials.

```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|inbound
|link:#inbound[inbound]
a|Inbound CHAP credentials.
```

```
|outbound
|link:#outbound[outbound]
a|Output CHAP credentials.
```

```
|===
```

```
[#ip_info]
[.api-collapsible-fifth-title]
ip_info
```

IP information

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|address
|string
a|IPv4 or IPv6 address
```

```
|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.
```

```
|===
```

```
[#ip_address_range]
[.api-collapsible-fifth-title]
ip_address_range
```

IP address range

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|end
|string
a|IPv4 or IPv6 address
```

```
|start
|string
a|IPv4 or IPv6 address
```

```
|===
```

```
[#initiator_address]
[.api-collapsible-fifth-title]
initiator_address
```

Initiator address ranges.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|masks
|array[link:#ip_info[ip_info]]
a|
```

```
|ranges
|array[link:#ip_address_range[ip_address_range]]
a|
```

|===

[#svm]

[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#iscsi_credentials]

[.api-collapsible-fifth-title]

iscsi_credentials

[cols=3*,options=header]

|===

|Name

|Type

|Description

|authentication_type

|string

a|The iSCSI authentication type. Required in POST and optional in PATCH.

|chap

|link:#chap[chap]

a|Challenge-Handshake Authentication Protocol (CHAP) credentials.

```

|initiator_address
|link:#initiator_address[initiator_address]
a|Initiator address ranges.

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]

```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
:leveloffset: -1
```

```
= Manage iSCSI services
```

```
:leveloffset: +1
```

```
[[ID191f5b695bfb5f763cef27b92bba0c22]]
```

```
= Protocols SAN iSCSI services endpoint overview
```

```
== Overview
```

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can log in to the SVM.

The iSCSI service REST API allows you to create, update, delete, and discover iSCSI services for SVMs.

== Performance monitoring

Performance of the SVM can be monitored by the ``metric.++`` and ``statistics.++`` properties. These show the performance of the SVM in terms of IOPS, latency and throughput. The ``metric.++`` properties denote an average whereas ``statistics.++`` properties denote a real-time monotonically increasing value aggregated across all nodes.

== Examples

=== Creating an iSCSI service for an SVM

The simplest way to create an iSCSI service is to specify only the SVM, either by name or UUID. By default, the new iSCSI service is enabled and uses the SVM name as its target alias.

In this example, the ``return_records`` query parameter is used to retrieve the new iSCSI service object in the REST response.

The API:

POST /api/protocols/san/iscsi/services

The call:

```
curl -X POST 'https://<mgmt-  
ip>/api/protocols/san/iscsi/services?return_records=true' -H 'accept:  
application/hal+json' -d '{ "svm": { "name": "svm1" } }'
```

The response:

```
{  
  "num_records": 1,  
  "records": [  
    {  
      "svm": {  
        "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",  
        "name": "svm1",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"  
          }  
        }  
      },  
      "enabled": true,  
      "target": {
```

```

    "name": "iqn.1992-
08.com.netapp:sn.19d04b8e94d711e88370005056b48fd2:vs.4",
    "alias": "svm1"
  },
  "_links": {
    "self": {
      "href": "/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-
8370-005056b48fd2"
    }
  }
}
]
}
-----
'''

```

=== Retrieving the iSCSI services for all SVMs in the cluster

```

-----

# The API:
GET /api/protocols/san/iscsi/services

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/iscsi/services' -H
'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-
8370-005056b48fd2"
        }
      }
    }
  ]
}

```



```

    }
  },
  {
    "svm": {
      "uuid": "25f617cf-94d7-11e8-8370-005056b48fd2",
      "name": "svm2",
      "_links": {
        "self": {
          "href": "/api/svm/svms/25f617cf-94d7-11e8-8370-005056b48fd2"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/services/25f617cf-94d7-11e8-8370-005056b48fd2"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/services"
  }
}
}
}
-----
'''

```

=== Retrieving details for a specific iSCSI service

The iSCSI service is identified by the UUID of its SVM.

The API:

GET /api/protocols/san/iscsi/services/{svm.uuid}

The call:

curl -X GET 'https://<mgmt-ip>/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-005056b48fd2' -H 'accept: application/hal+json'

The response:

{

```

"svm": {
  "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
  "name": "svm1",
  "_links": {
    "self": {
      "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
    }
  }
},
"enabled": true,
"target": {
  "name": "iqn.1992-
08.com.netapp:sn.19d04b8e94d711e88370005056b48fd2:vs.4",
  "alias": "svm1"
},
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-
005056b48fd2"
  }
}
}
}

```

'''

=== Disabling an iSCSI service

Disabling an iSCSI service shuts down all active iSCSI sessions for the SVM and prevents the creation of new iSCSI sessions.

The iSCSI service to update is identified by the UUID of its SVM.

The API:

PATCH /api/protocols/san/iscsi/services/{svm.uuid}

The call:

```

curl -X PATCH 'https://<mgmt-
ip>/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-005056b48fd2'
-H 'accept: application/hal+json' -d '{ "enabled": "false" }'

```

You can retrieve the iSCSI service to confirm the change.

In this example, the `fields` query parameter is used to limit the response to the `enabled` property and iSCSI service identifiers.

The API:

```
GET /api/protocols/san/iscsi/services/{svm.uuid}
```

The call:

```
curl -X GET 'https://<mgmt-ip>/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-005056b48fd2?fields=enabled' -H 'accept: application/hal+json'
```

The response:

```
{
  "svm": {
    "uuid": "19d04b8e-94d7-11e8-8370-005056b48fd2",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/19d04b8e-94d7-11e8-8370-005056b48fd2"
      }
    }
  },
  "enabled": false,
  "_links": {
    "self": {
      "href": "/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-005056b48fd2"
    }
  }
}
```

'''

=== Deleting an iSCSI service

The iSCSI service must be disabled before it can be deleted.

The iSCSI service to be deleted is identified by the UUID of its SVM.

The API:

```
DELETE /api/protocols/san/iscsi/services/{svm.uuid}
```

```

# The call:
curl -X DELETE 'https://<mgmt-
ip>/api/protocols/san/iscsi/services/19d04b8e-94d7-11e8-8370-005056b48fd2'
-H 'accept: application/hal+json'
----

[[IDf23b292db681a49c7675fd8c8fe48112]]
= Retrieve iSCSI services

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/san/iscsi/services`#

*Introduced In:* 9.6

Retrieves iSCSI services.

== Expensive properties

There is an added cost to retrieving values for these properties. They are
not included by default in GET results and must be explicitly requested
using the `fields` query parameter. See
xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin
g_specific_fields[Requesting specific fields] to learn more.

* `statistics.+++`
* `metric.+++`

== Related ONTAP commands

* `vserver iscsi show`

== Learn more

*
xref:{relative_path}protocols_san_iscsi_services_endpoint_overview.html[DO
C /protocols/san/iscsi/services]

== Parameters

[cols=5*,options=header]
|===

```

```
|Name
|Type
|In
|Required
|Description

|metric.throughput.write
|integer
|query
|False
a|Filter by metric.throughput.write
```

* Introduced in: 9.7

```
|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read
```

* Introduced in: 9.7

```
|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total
```

* Introduced in: 9.7

```
|metric.duration
|string
|query
|False
a|Filter by metric.duration
```

* Introduced in: 9.7

```
|metric.timestamp
|string
|query
|False
a|Filter by metric.timestamp
```

* Introduced in: 9.7

```
|metric.status
|string
|query
|False
a|Filter by metric.status
```

* Introduced in: 9.7

```
|metric.iops.total
|integer
|query
|False
a|Filter by metric.iops.total
```

* Introduced in: 9.7

```
|metric.iops.read
|integer
|query
|False
a|Filter by metric.iops.read
```

* Introduced in: 9.7

```
|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other
```

* Introduced in: 9.7

```
|metric.iops.write
|integer
|query
|False
a|Filter by metric.iops.write
```

* Introduced in: 9.7

```
|metric.latency.total
|integer
|query
|False
a|Filter by metric.latency.total
```

* Introduced in: 9.7

```
|metric.latency.read
|integer
|query
|False
a|Filter by metric.latency.read
```

* Introduced in: 9.7

```
|metric.latency.other
|integer
|query
|False
a|Filter by metric.latency.other
```

* Introduced in: 9.7

```
|metric.latency.write
|integer
|query
|False
a|Filter by metric.latency.write
```

* Introduced in: 9.7

```
|svm.uuid
|string
|query
|False
a|Filter by svm.uuid
```

```
|svm.name
|string
```

```
|query
|False
a|Filter by svm.name

|enabled
|boolean
|query
|False
a|Filter by enabled

|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total

* Introduced in: 9.7

|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read

* Introduced in: 9.7

|statistics.latency_raw.other
|integer
|query
|False
a|Filter by statistics.latency_raw.other

* Introduced in: 9.7

|statistics.latency_raw.write
|integer
|query
|False
a|Filter by statistics.latency_raw.write

* Introduced in: 9.7
```



```
|statistics.timestamp  
|string  
|query  
|False  
a|Filter by statistics.timestamp
```

* Introduced in: 9.7

```
|statistics.iops_raw.total  
|integer  
|query  
|False  
a|Filter by statistics.iops_raw.total
```

* Introduced in: 9.7

```
|statistics.iops_raw.read  
|integer  
|query  
|False  
a|Filter by statistics.iops_raw.read
```

* Introduced in: 9.7

```
|statistics.iops_raw.other  
|integer  
|query  
|False  
a|Filter by statistics.iops_raw.other
```

* Introduced in: 9.7

```
|statistics.iops_raw.write  
|integer  
|query  
|False  
a|Filter by statistics.iops_raw.write
```

* Introduced in: 9.7

```
|statistics.throughput_raw.write
```

```
|integer
|query
|False
a|Filter by statistics.throughput_raw.write
```

* Introduced in: 9.7

```
|statistics.throughput_raw.read
|integer
|query
|False
a|Filter by statistics.throughput_raw.read
```

* Introduced in: 9.7

```
|statistics.throughput_raw.total
|integer
|query
|False
a|Filter by statistics.throughput_raw.total
```

* Introduced in: 9.7

```
|statistics.status
|string
|query
|False
a|Filter by statistics.status
```

* Introduced in: 9.7

```
|target.alias
|string
|query
|False
a|Filter by target.alias
```

```
|target.name
|string
|query
|False
a|Filter by target.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

```
|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records.

|records
|array[link:#iscsi_service[iscsi_service]]
a|

|===
```

```
.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
```

```

    }
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },

```

```

    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "target": {
      "alias": "svm1",
      "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",

```

```

        "message": "string"
    }
],
"code": "4",
"message": "entry doesn't exist",
"target": "uuid"
}
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]

```

```

a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|self
|link:#href[href]
a|

```

```

|===

```

```

[#iops]
[.api-collapsible-fifth-title]
iops

```

The rate of I/O operations observed at the storage object.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|other
|integer

```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```

|read
|integer

```



```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|read
|integer

a|Performance metric for read I/O operations.

|total
|integer

a|Performance metric aggregated over all types of I/O operations.

|write
|integer

a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

[cols=3*,options=header]

|===

|Name
|Type
|Description

```

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
```

```
[.api-collapsible-fifth-title]
```

```
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency

per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#statistics]
[.api-collapsible-fifth-title]
statistics

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be

```

returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
```

a|The unique identifier of the SVM.

|===

[#target]
[.api-collapsible-fifth-title]
target

[cols=3*,options=header]

|===

|Name
|Type
|Description

|alias
|string

a|The iSCSI target alias of the iSCSI service.

The target alias can contain one (1) to 128 characters and feature any printable character except space (" "). A PATCH request with an empty alias ("") clears the alias.

Optional in POST and PATCH. In POST, this defaults to the name of the SVM.

|name
|string

a|The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.

If required, the target name can be modified using the ONTAP command line.

- * example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2
- * maxLength: 128
- * minLength: 1
- * readOnly: 1
- * Introduced in: 9.6

|===

[#iscsi_service]
[.api-collapsible-fifth-title]
iscsi_service

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can log in to the SVM.

An iSCSI service is identified by the UUID of its SVM.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|enabled
```

```
|boolean
```

a|The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.

Optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
```

```
|link:#metric[metric]
```

```
a|
```

```
|statistics
```

```
|link:#statistics[statistics]
```

```
a|
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|target
```

```
|link:#target[target]
```

```
a|
```

```
|===
```

```
[#error_arguments]
```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

```

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

[[IDcf8a767656efb8ea7736449d7530d7b2]]
= Create an iSCSI service

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/san/iscsi/services`#

*Introduced In:* 9.6

Creates an iSCSI service.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the iSCSI
service.

== Related ONTAP commands

* `vserver iscsi create`

== Learn more

*
xref:{relative_path}protocols_san_iscsi_services_endpoint_overview.html[DO
C /protocols/san/iscsi/services]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In

```

```

|Required
|Description

|return_records
|boolean
|query
|False
a|The default is false.  If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|The administrative state of the iSCSI service. The iSCSI service can be
disabled to block all iSCSI connectivity to the SVM.

Optional in POST and PATCH. The default setting is _true_ (enabled) in
POST.

|metric
|link:#metric[metric]
a|

|statistics
|link:#statistics[statistics]
a|

|svm
|link:#svm[svm]
a|

|target
|link:#target[target]
a|

```

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "metric": {
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    }
  },
}
```

```

    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "alias": "svm1",
    "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
  }
}
====

```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#iscsi_service[iscsi_service]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "metric": {
        "duration": "PT15S",
        "iops": {

```

```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"statistics": {
    "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"target": {
    "alias": "svm1",
    "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
}
}

```

```
]
}
====

== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error Code | Description

| 1115127
| The cluster lacks a valid iSCSI license.

| 2621462
| The supplied SVM does not exist.

| 2621507
| The iSCSI protocol is not allowed for the specified SVM.

| 2621706
| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.

| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5373966
| An iSCSI service cannot be created in an SVM that is configured for
NVMe.

| 5374077
| An iSCSI service already exists for the specified SVM.

| 5374893
| The SVM is stopped. The SVM must be running to create an iSCSI service.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```



```
====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#latency]

[.api-collapsible-fifth-title]

latency

The round trip latency in microseconds observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw

The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```

|write
|integer
a|Performance metric for write I/O operations.

|===

[#statistics]
[.api-collapsible-fifth-title]
statistics

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

```

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#target]
[.api-collapsible-fifth-title]
target
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```


|Type
|Description

|alias
|string

a|The iSCSI target alias of the iSCSI service.

The target alias can contain one (1) to 128 characters and feature any printable character except space (" "). A PATCH request with an empty alias ("") clears the alias.

Optional in POST and PATCH. In POST, this defaults to the name of the SVM.

|name
|string

a|The iSCSI target name of the iSCSI service. This is generated for the SVM during POST.

If required, the target name can be modified using the ONTAP command line.

- * example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2
- * maxLength: 128
- * minLength: 1
- * readOnly: 1
- * Introduced in: 9.6

|===

[#iscsi_service]
[.api-collapsible-fifth-title]
iscsi_service

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can log in to the SVM.

An iSCSI service is identified by the UUID of its SVM.

[cols=3*,options=header]
|===

|Name
|Type
|Description

|enabled
|boolean
a|The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.

Optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

|metric
|link:#metric[metric]
a|

|statistics
|link:#statistics[statistics]
a|

|svm
|link:#svm[svm]
a|

|target
|link:#target[target]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID2d882f488890187ea877c51d2306c370]]
= Delete an iSCSI service

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
```

```
block]#`/protocols/san/iscsi/services/{svm.uuid}`#
```

Introduced In: 9.6

Deletes an iSCSI service. An iSCSI service must be disabled before it can be deleted.

== Related ONTAP commands

* `vserver iscsi delete`

== Learn more

*

xref:{relative_path}protocols_san_iscsi_services_endpoint_overview.html[DOC /protocols/san/iscsi/services]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|The unique identifier of the SVM for which to delete the iSCSI service.

|===

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 2621462

| An SVM with the specified UUID does not exist.

| 5373960

| The iSCSI service is enabled. The iSCSI service must be disabled before it can be deleted.

| 5374078

| The SVM does not have an iSCSI service.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

```

}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[ID03191ff364f05fe0f845194adef6859f]]
= Retrieve an iSCSI service
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/san/iscsi/services/{svm.uuid}`#
```

```
*Introduced In:* 9.6
```

```
Retrieves an iSCSI service.
```

```
== Related ONTAP commands
```

```
* `vserver iscsi show`
```

```
== Learn more
```

```
*
```

```
xref:{relative_path}protocols_san_iscsi_services_endpoint_overview.html[DO
C /protocols/san/iscsi/services]
```

```
== Parameters
```

```

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|The unique identifier of the SVM for which to retrieve the iSCSI
service.

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|The administrative state of the iSCSI service. The iSCSI service can be
disabled to block all iSCSI connectivity to the SVM.

Optional in POST and PATCH. The default setting is _true_ (enabled) in
POST.

```



```

|metric
|link:#metric[metric]
a|

|statistics
|link:#statistics[statistics]
a|

|svm
|link:#svm[svm]
a|

|target
|link:#target[target]
a|

|===

```

```

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",

```

```

    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"statistics": {
  "iops_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"target": {
  "alias": "svm1",
  "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
}
}
=====

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 2621462
| An SVM with the specified UUID does not exist.
| 5374078
| The SVM does not have an iSCSI service.
|===

```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|error
```

```
|link:#error[error]
```

```
a|
```

```
|===
```

.Example error

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

```

    }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#iops]
[.api-collapsible-fifth-title]

```

iops

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be  
metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```

|write
|integer
a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with

```

"backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput

|link:#throughput[throughput]

a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#iops_raw]

[.api-collapsible-fifth-title]

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.


```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]  
[.api-collapsible-fifth-title]  
throughput_raw
```

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|read  
|integer  
a|Performance metric for read I/O operations.
```

```
|total  
|integer  
a|Performance metric aggregated over all types of I/O operations.
```

```
|write  
|integer  
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]  
[.api-collapsible-fifth-title]  
statistics
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#svm]
[.api-collapsible-fifth-title]

svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#target]

[.api-collapsible-fifth-title]

target

[cols=3*,options=header]

|===

|Name

|Type

|Description

|alias

|string

a|The iSCSI target alias of the iSCSI service.

The target alias can contain one (1) to 128 characters and feature any printable character except space (" "). A PATCH request with an empty alias ("") clears the alias.

Optional in POST and PATCH. In POST, this defaults to the name of the SVM.

|name

```
|string
a|The iSCSI target name of the iSCSI service. This is generated for the
SVM during POST.
```

If required, the target name can be modified using the ONTAP command line.

```
* example: ign.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2
* maxLength: 128
* minLength: 1
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
```

```

|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

[[ID99e21b24a1f64095762666db7618330e]]
= Update an iSCSI service

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/protocols/san/iscsi/services/{svm.uuid}`#

*Introduced In:* 9.6

Updates an iSCSI service.

== Related ONTAP commands

* `vserver iscsi modify`
* `vserver iscsi start`
* `vserver iscsi stop`

== Learn more

```

*
xref:{relative_path}protocols_san_iscsi_services_endpoint_overview.html[DO
C /protocols/san/iscsi/services]

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|svm.uuid
|string
|path
|True
a|The unique identifier of the SVM for which to update the iSCSI service.

|===

== Request Body

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.

Optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

|metric
|link:#metric[metric]
a|

```
|statistics
|link:#statistics[statistics]
a|
```

```
|svm
|link:#svm[svm]
a|
```

```
|target
|link:#target[target]
a|
```

```
|===
```

```
.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
```

```
{
  "metric": {
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    }
  }
}
```



```

    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "target": {
    "alias": "svm1",
    "name": "iqn.1992-
08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2"
  }
}
====

== Response

```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```

|====
| Error Code | Description
|
| 2621462
| An SVM with the specified UUID does not exist.
|
| 5374078
| The SVM does not have an iSCSI service.
|====

```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
```

====

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|href
|string
a|
```

|===

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```

|Type
|Description

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized
failure. Whenever a sample collection is missed but done at a later time,
it is back filled to the previous 15 second timestamp and tagged with
"backfilled_data". "Inconsistent_delta_time" is encountered when the time
between two collections is not the same for all nodes. Therefore, the
aggregated value might be over or under inflated. "Negative_delta" is
returned when an expected monotonically increasing value has decreased in
value. "Inconsistent_old_data" is returned when one or more nodes do not
have the latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

```

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
```

a|Performance metric for read I/O operations.

```
|total
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This should be divided by the raw IOPS value to calculate the average latency per I/O operation.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw

Throughput bytes observed at the storage object. This should be used along
with delta time to calculate the rate of throughput bytes per unit of
time.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#statistics]
[.api-collapsible-fifth-title]
statistics

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This should
be used along with delta time to calculate the rate of I/O operations per
unit of time.

|latency_raw
|link:#latency_raw[latency_raw]
a|The raw latency in microseconds observed at the storage object. This
should be divided by the raw IOPS value to calculate the average latency
per I/O operation.

|status
|string
a|Any errors associated with the sample. For example, if the aggregation
of data over multiple nodes fails then any of the partial errors might be
returned, "ok" on success, or "error" on any internal uncategorized

```


failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This should be used
along with delta time to calculate the rate of throughput bytes per unit
of time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#target]
[.api-collapsible-fifth-title]
target
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|alias
|string
a|The iSCSI target alias of the iSCSI service.
```

The target alias can contain one (1) to 128 characters and feature any printable character except space (" "). A PATCH request with an empty alias ("") clears the alias.

Optional in POST and PATCH. In POST, this defaults to the name of the SVM.

```
|name
|string
a|The iSCSI target name of the iSCSI service. This is generated for the
SVM during POST.
```

If required, the target name can be modified using the ONTAP command line.

```
* example: iqn.1992-08.com.netapp:sn.574caf71890911e8a6b7005056b4ea79:vs.2
* maxLength: 128
* minLength: 1
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#iscsi_service]
[.api-collapsible-fifth-title]
iscsi_service
```

An iSCSI service defines the properties of the iSCSI target for an SVM. There can be at most one iSCSI service for an SVM. An SVM's iSCSI service must be created before iSCSI initiators can log in to the SVM.

An iSCSI service is identified by the UUID of its SVM.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

a|The administrative state of the iSCSI service. The iSCSI service can be disabled to block all iSCSI connectivity to the SVM.

Optional in POST and PATCH. The default setting is `_true_` (enabled) in POST.

```
|metric
```

```
|link:#metric[metric]
```

```
a|
```

```
|statistics
```

```
|link:#statistics[statistics]
```

```
a|
```

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|target
```

```
|link:#target[target]
```

```
a|
```

```
|===
```

```
[#error_arguments]
```

```
[.api-collapsible-fifth-title]
```

```
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

```

====

[[ID0839885c41a9821109d5fa11eb3a7df7]]

= Retrieve historical performance metrics for the iSCSI protocol of an SVM

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/san/iscsi/services/{svm.uuid}/metrics`#

Introduced In: 9.7

Retrieves historical performance metrics for the iSCSI protocol of an SVM.

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|status

|string

|query

|False

a|Filter by status

|iops.total

|integer

|query

|False

a|Filter by iops.total

|iops.read

|integer

|query

|False

a|Filter by iops.read

```
|iops.other
|integer
|query
|False
a|Filter by iops.other
```

```
|iops.write
|integer
|query
|False
a|Filter by iops.write
```

```
|latency.total
|integer
|query
|False
a|Filter by latency.total
```

```
|latency.read
|integer
|query
|False
a|Filter by latency.read
```

```
|latency.other
|integer
|query
|False
a|Filter by latency.other
```

```
|latency.write
|integer
|query
|False
a|Filter by latency.write
```

```
|throughput.total
|integer
|query
|False
a|Filter by throughput.total
```

```
|throughput.read  
|integer  
|query  
|False  
a|Filter by throughput.read
```

```
|throughput.other  
|integer  
|query  
|False  
a|Filter by throughput.other
```

```
|throughput.write  
|integer  
|query  
|False  
a|Filter by throughput.write
```

```
|duration  
|string  
|query  
|False  
a|Filter by duration
```

```
|timestamp  
|string  
|query  
|False  
a|Filter by timestamp
```

```
|svm.uuid  
|string  
|path  
|True  
a|The unique identifier of the SVM.
```

```
|interval  
|string  
|query
```

```

|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:

* 1h: Metrics over the most recent hour sampled over 15 seconds.
* 1d: Metrics over the most recent day sampled over 5 minutes.
* 1w: Metrics over the most recent week sampled over 30 minutes.
* 1m: Metrics over the most recent month sampled over 2 hours.
* 1y: Metrics over the most recent year sampled over a day.
* Default value: 1
* enum: ["1h", "1d", "1w", "1m", "1y"]

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

```



```

|return_records
|boolean
|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#records[records]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {

```

```

    "href": "/api/resourcelink"
  },
  "self": {
    "href": "/api/resourcelink"
  }
},
"records": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  }
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|error
|link:#error[error]
a|
```

```
|===
```

```
.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====
```

```
== Definitions
```

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
```

```

a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#iops]
[.api-collapsible-fifth-title]
iops

```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

```
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#records]
[.api-collapsible-fifth-title]
records
```

Performance numbers, such as IOPS latency and throughput.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```



```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= View iSCSI sessions

```

```
:leveloffset: +1
```

```
[[ID9f2feddc450d80924f4d9b98ceb08f20]]  
= Protocols SAN iSCSI sessions endpoint overview
```

== Overview

An iSCSI session is one or more TCP connections that link an iSCSI initiator with an iSCSI target. TCP connections can be added and removed from an iSCSI session by the iSCSI initiator. Across all TCP connections within an iSCSI session, an initiator sees one and the same target. After the connection is established, iSCSI control, data, and status messages are communicated over the session.

The iSCSI sessions REST API provides information about iSCSI initiators that have successfully logged in to ONTAP.

== Examples

=== Retrieving all iSCSI sessions

```
----
```

```
# The API:
```

```
GET /api/protocols/san/iscsi/sessions
```

```
# The call:
```

```
curl -X GET "https://<mgmt-ip>/api/protocols/san/iscsi/sessions" -H  
"accept: application/hal+json"
```

```
# The response:
```

```
{  
  "records": [  
    {  
      "svm": {  
        "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",  
        "name": "svm1",  
        "_links": {  
          "self": {  
            "href": "/api/svm/svms/a009a9e7-4081-b576-7575-ada21efcaf16"  
          }  
        }  
      },  
      "target_portal_group": "iscsi_lif1",
```

```

    "tsih": 10,
    "_links": {
      "self": {
        "href": "/api/protocols/san/iscsi/sessions/a009a9e7-4081-b576-7575-ada21efcaf16/iscsi_lif1/10"
      }
    },
    {
      "svm": {
        "uuid": "b009a9e7-4081-b576-7575-ada21efcaf16",
        "name": "svm2",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b009a9e7-4081-b576-7575-ada21efcaf16"
          }
        }
      },
      "target_portal_group": "iscsi_lif2",
      "tsih": 11,
      "_links": {
        "self": {
          "href": "/api/protocols/san/iscsi/sessions/b009a9e7-4081-b576-7575-ada21efcaf16/iscsi_lif2/11"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/protocols/san/iscsi/sessions"
    }
  }
}
----

'''

```

=== Retrieving all of the iSCSI sessions under the target portal group `_iscsi_lif1_`

The ``tpgroup`` query parameter is used to perform the query.

```

# The API:
GET /api/protocols/san/iscsi/sessions

# The call:
curl -X GET "https://<mgmt-
ip>/api/protocols/san/iscsi/sessions?tpgroup=iscsi_lif1" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/a009a9e7-4081-b576-7575-ada21efcaf16"
          }
        }
      },
      "target_portal_group": "iscsi_lif1",
      "tsih": 10,
      "_links": {
        "self": {
          "href": "/api/protocols/san/iscsi/sessions/a009a9e7-4081-b576-
7575-ada21efcaf16/iscsi_lif1/10"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/protocols/san/iscsi/sessions"
    }
  }
}
----

'''

=== Retrieving an iSCSI session

----

```

```

# The API:
GET
/api/protocols/san/iscsi/sessions/{svm.uuid}/{target_portal_group}/{tsih}

# The call:
curl -X GET "https://<mgmt-ip>/api/protocols/san/iscsi/sessions/a009a9e7-4081-b576-7575-ada21efcaf16/iscsi_lif1/10" -H "accept: application/hal+json"

# The response:
{
  "svm": {
    "uuid": "a009a9e7-4081-b576-7575-ada21efcaf16",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/a009a9e7-4081-b576-7575-ada21efcaf16"
      }
    }
  },
  "target_portal_group": "iscsi_lif1",
  "tsih": 10,
  "initiator": {
    "name": "iqn.1994-05.com.example:string"
  },
  "isid": "61:62:63:64:65:00",
  "target_portal_group_tag": 1027,
  "connections": [
    {
      "cid": 1,
      "authentication_type": "chap",
      "initiator_address": {
        "address": "10.224.123.85",
        "port": 43827
      },
      "interface": {
        "name": "iscsi_lif1",
        "uuid": "c15439b4-dbb4-11e8-90ac-005056bba882",
        "ip": {
          "address": "192.168.0.1",
          "port": 3260
        },
        "_links": {
          "self": {
            "href": "/api/network/ip/interfaces/c15439b4-dbb4-11e8-90ac-005056bba882"
          }
        }
      }
    }
  ]
}

```

```

    }
  }
}
],
"igroups": [
  {
    "uuid": "af7838cd-f993-4faf-90b7-5524787ae1e8",
    "name": "igroup1",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/af7838cd-f993-4faf-90b7-5524787ae1e8"
      }
    }
  },
  {
    "uuid": "bf7838cd-f993-4faf-90b7-5524787ae1e8",
    "name": "igroup2",
    "_links": {
      "self": {
        "href": "/api/protocols/san/igroups/bf7838cd-f993-4faf-90b7-5524787ae1e8"
      }
    }
  }
],
"_links": {
  "self": {
    "href": "/api/protocols/san/iscsi/sessions/a009a9e7-4081-b576-7575-ada21efcaf16/iscsi_lif1/10"
  }
}
}
}
-----

```

```
[[ID650c9f32ec2d97cf9b4567b3ce6028d1]]
```

```
= Retrieve iSCSI sessions
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/san/iscsi/sessions`#
```

Introduced In: 9.6

Retrieves iSCSI sessions.

== Related ONTAP commands

- * ``vserver iscsi connection show``
- * ``vserver iscsi session parameter show``
- * ``vserver iscsi session show``

== Learn more

*

xref:{relative_path}protocols_san_iscsi_sessions_endpoint_overview.html[DO
C /protocols/san/iscsi/sessions]

== Parameters

[cols=5*,options=header]
|==

|Name
|Type
|In
|Required
|Description

|tsih
|integer
|query
|False
a|Filter by tsih

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

```
|initiator.name  
|string  
|query  
|False  
a|Filter by initiator.name
```

```
|initiator.alias  
|string  
|query  
|False  
a|Filter by initiator.alias
```

```
|isid  
|string  
|query  
|False  
a|Filter by isid
```

```
|connections.initiator_address.port  
|integer  
|query  
|False  
a|Filter by connections.initiator_address.port
```

```
|connections.initiator_address.address  
|string  
|query  
|False  
a|Filter by connections.initiator_address.address
```

```
|connections.authentication_type  
|string  
|query  
|False  
a|Filter by connections.authentication_type
```

```
|connections.cid  
|integer  
|query
```



```
|False
a|Filter by connections.cid

|connections.interface.uuid
|string
|query
|False
a|Filter by connections.interface.uuid

|connections.interface.name
|string
|query
|False
a|Filter by connections.interface.name

|connections.interface.ip.address
|string
|query
|False
a|Filter by connections.interface.ip.address

|connections.interface.ip.port
|integer
|query
|False
a|Filter by connections.interface.ip.port

|target_portal_group_tag
|integer
|query
|False
a|Filter by target_portal_group_tag

|target_portal_group
|string
|query
|False
a|Filter by target_portal_group

|igroups.uuid
```

```
|string
|query
|False
a|Filter by igroups.uuid
```

```
|igroups.name
|string
|query
|False
a|Filter by igroups.name
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.
```

* Default value: 1

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of records.

|records

|array[link:#iscsi_session[iscsi_session]]

a|

|===

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

{

 "_links": {

 "next": {

```

    "href": "/api/resourcelink"
  },
  "self": {
    "href": "/api/resourcelink"
  }
},
"records": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "connections": [
      {
        "_links": {
          "next": {
            "href": "/api/resourcelink"
          },
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "authentication_type": "string",
        "cid": 0,
        "initiator_address": {
          "address": "10.10.10.7",
          "port": "55432"
        },
        "interface": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "ip": {
            "address": "10.10.10.7",
            "port": "3260"
          },
          "name": "lif1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      }
    ],
    "igroups": [
      {

```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "igroup1",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  }
],
"initiator": {
  "alias": "initiator_alias1",
  "name": "iqn.1992-01.example.com:string"
},
"isid": "61:62:63:64:65:00",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"target_portal_group": "tpgroup1",
"target_portal_group_tag": 0,
"tsih": 0
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

```

```

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|next
|link:href[href]
a|
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:href[href]
a|
```

```
|===
```

```
[#initiator_address]
[.api-collapsible-fifth-title]
initiator_address
```

The TCP socket information for the initiator end of the connection. This is useful for network packet debugging.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|address
|string
a|The TCP IPv4 or IPv6 address of the initiator end of the iSCSI
connection.

|port
|integer
a|The TCP port number of the initiator end of the iSCSI connection.

```

```

|===

```

```

[#ip]
[.api-collapsible-fifth-title]
ip

```

The IP information. ONTAP only supports port 3260.

```

[cols=3*,options=header]

```

```

|===

```

```

|Name
|Type
|Description

```

```

|address
|string
a|IPv4 or IPv6 address

```

```

|port
|integer
a|The TCP port number of the iSCSI access endpoint.

```

```

|===

```

```

[#interface]
[.api-collapsible-fifth-title]

```



```
interface
```

The network interface information for the target end of the connection.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|ip
```

```
|link:#ip[ip]
```

```
a|The IP information. ONTAP only supports port 3260.
```

```
* readOnly: 1
```

```
* Introduced in: 9.6
```

```
|name
```

```
|string
```

```
a|The name of the interface.
```

```
|uuid
```

```
|string
```

```
a|The UUID that uniquely identifies the interface.
```

```
|===
```

```
[#iscsi_connection]
```

```
[.api-collapsible-fifth-title]
```

```
iscsi_connection
```

An active iSCSI connection.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```

|Description

|_links
|link:#_links[_links]
a|

|authentication_type
|string
a|The iSCSI authentication type used to establish the connection.

|cid
|integer
a|The identifier of the connection within the session.

|initiator_address
|link:#initiator_address[initiator_address]
a|The TCP socket information for the initiator end of the connection. This
is useful for network packet debugging.

|interface
|link:#interface[interface]
a|The network interface information for the target end of the connection.

|===

[#igroups]
[.api-collapsible-fifth-title]
igroups

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the initiator group.

```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#initiator]
[.api-collapsible-fifth-title]
initiator
```

The initiator that created the session.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|alias
```

```
|string
```

```
a|The initiator alias.
```

```
|name
```

```
|string
```

```
a|The world wide unique name of the initiator.
```

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the SVM.
```

```
|===
```

```
[#iscsi_session]
```

```
[.api-collapsible-fifth-title]
```

```
iscsi_session
```

An iSCSI session is one or more TCP connections that link an iSCSI initiator with an iSCSI target. TCP connections can be added and removed from an iSCSI session by the iSCSI initiator. Across all TCP connections within an iSCSI session, an initiator sees one and the same target. After the connection is established, iSCSI control, data, and status messages are communicated over the session.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|connections
```

```
|array[link:#iscsi_connection[iscsi_connection]]
```

```
a|The iSCSI connections that make up the iSCSI session.
```

```
|igroups
```

```
|array[link:#igroups[igroups]]
```

```
a|The initiator groups in which the initiator is a member.
```

```

|initiator
|link:#initiator[initiator]
a|The initiator that created the session.

|isid
|string
a|The initiator portion of the session identifier specified by the
initiator during login.

|svm
|link:#svm[svm]
a|

|target_portal_group
|string
a|The target portal group to which the session belongs.

|target_portal_group_tag
|integer
a|The target portal group tag of the session.

|tsih
|integer
a|The target session identifier handle (TSIH) of the session.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

[[ID8bc877696ccd8c17eb9655c2c3dac86f]]

= Retrieve an iSCSI session

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/protocols/san/iscsi/sessions/{svm.uuid}/{tpgroup}/{tsih}`#

Introduced In: 9.6

Retrieves an iSCSI session.

== Related ONTAP commands

* `vserver iscsi connection show`
* `vserver iscsi session parameter show`
* `vserver iscsi session show`

== Learn more

*

xref:{relative_path}protocols_san_iscsi_sessions_endpoint_overview.html[DOC /protocols/san/iscsi/sessions]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|svm.uuid

|string

|path

|True

a|The unique identifier of the SVM of the iSCSI session.

|tpgroup

|string

|path

|True

a|The target portal group of the iSCSI session.

```
|tsih
|integer
|path
|True
a|The target session identifying handle.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|connections
|array[link:#iscsi_connection[iscsi_connection]]
a|The iSCSI connections that make up the iSCSI session.

|igroups
|array[link:#igroups[igroups]]
a|The initiator groups in which the initiator is a member.

|initiator
|link:#initiator[initiator]
a|The initiator that created the session.
```



```
|isid
|string
a|The initiator portion of the session identifier specified by the
initiator during login.
```

```
|svm
|link:#svm[svm]
a|
```

```
|target_portal_group
|string
a|The target portal group to which the session belongs.
```

```
|target_portal_group_tag
|integer
a|The target portal group tag of the session.
```

```
|tsih
|integer
a|The target session identifier handle (TSIH) of the session.
```

```
|===
```

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "connections": [
    {
      "_links": {
        "next": {
          "href": "/api/resourcelink"
        },
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  ]
}
```

```

    },
    "authentication_type": "string",
    "cid": 0,
    "initiator_address": {
      "address": "10.10.10.7",
      "port": "55432"
    },
    "interface": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "ip": {
        "address": "10.10.10.7",
        "port": "3260"
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "igroups": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "igroup1",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "initiator": {
    "alias": "initiator_alias1",
    "name": "iqn.1992-01.example.com:string"
  },
  "isid": "61:62:63:64:65:00",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}

```

```

    },
    "target_portal_group": "tpgroup1",
    "target_portal_group_tag": 0,
    "tsih": 0
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
|
| 2621462
| An SVM with the specified UUID does not exist.
|===

```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```
|===
```

.Example error

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ]
  }
}

```

```

    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

```

```
[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|next
|link:#href[href]
a|
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#initiator_address]
[.api-collapsible-fifth-title]
initiator_address
```

The TCP socket information for the initiator end of the connection. This is useful for network packet debugging.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|address
|string
a|The TCP IPv4 or IPv6 address of the initiator end of the iSCSI connection.
```

```
|port
|integer
a|The TCP port number of the initiator end of the iSCSI connection.
```

|===

[#ip]

[.api-collapsible-fifth-title]

ip

The IP information. ONTAP only supports port 3260.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|address

|string

a|IPv4 or IPv6 address

|port

|integer

a|The TCP port number of the iSCSI access endpoint.

|===

[#interface]

[.api-collapsible-fifth-title]

interface

The network interface information for the target end of the connection.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

```
|ip
|link:#ip[ip]
a|The IP information. ONTAP only supports port 3260.

* readOnly: 1
* Introduced in: 9.6
```

```
|name
|string
a|The name of the interface.
```

```
|uuid
|string
a|The UUID that uniquely identifies the interface.
```

```
|===
```

```
[#iscsi_connection]
[.api-collapsible-fifth-title]
iscsi_connection
```

An active iSCSI connection.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|authentication_type
|string
a|The iSCSI authentication type used to establish the connection.
```

```
|cid
|integer
a|The identifier of the connection within the session.
```

```
|initiator_address
|link:#initiator_address[initiator_address]
a|The TCP socket information for the initiator end of the connection. This
is useful for network packet debugging.
```

```
|interface
|link:#interface[interface]
a|The network interface information for the target end of the connection.
```

```
|===
```

```
[#igroups]
[.api-collapsible-fifth-title]
igroups
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the initiator group.
```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#initiator]
[.api-collapsible-fifth-title]
initiator
```

The initiator that created the session.


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|alias
|string
a|The initiator alias.


|name
|string
a|The world wide unique name of the initiator.


|===


[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.


|uuid
|string
a|The unique identifier of the SVM.


|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
= Manage SAN LUN maps
```

```
:leveloffset: +1
```

```
[[IDd63792c9ccfde6f752e52da511dbd45a]]
= Protocols SAN lun-maps endpoint overview
```

```
== Overview
```

A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between an initiator group and a LUN is many initiator groups to many LUNs.

The LUN map REST API allows you to create, delete, and discover LUN maps.

```
== Examples
```

```
=== Creating a LUN map
```

```
----
```

```
# The API:
POST /api/protocols/san/lun-maps
```

```
# The call:
```

```

curl -X POST 'https://<mgmt-ip>/api/protocols/san/lun-maps' -H 'accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "igroup": { "name":
"igroup1" }, "lun": { "name": "/vol/vol1/lun1" } }'
----

'''

=== Retrieving all of the LUN maps

----

# The API:
GET /api/protocols/san/lun-maps

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/lun-maps' -H 'accept:
application/hal+json'

# The response:
{
"records": [
  {
    "svm": {
      "uuid": "03157e81-24c5-11e9-9ec1-005056bba643",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/03157e81-24c5-11e9-9ec1-005056bba643"
        }
      }
    },
    "lun": {
      "uuid": "a60d9862-9bee-49a6-8162-20d2421bb1a6",
      "name": "/vol/vol1/lun1",
      "_links": {
        "self": {
          "href": "/api/storage/luns/a60d9862-9bee-49a6-8162-20d2421bb1a6"
        }
      }
    },
    "igroup": {
      "uuid": "40d98b2c-24c5-11e9-9ec1-005056bba643",
      "name": "ig1",
      "_links": {
        "self": {
          "href": "/api/protocols/san/igroups/40d98b2c-24c5-11e9-9ec1-

```

```

005056bba643"
    }
  }
},
"_links": {
  "self": {
    "href": "/api/protocols/san/lun-maps/a60d9862-9bee-49a6-8162-
20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643"
  }
}
],
"num_records": 1,
"_links": {
  "self": {
    "href": "/api/protocols/san/lun-maps"
  }
}
}
}
----

'''

=== Retrieving a specific LUN map

----

# The API:
GET /api/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/protocols/san/lun-maps/a60d9862-9bee-
49a6-8162-20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643' -H 'accept:
application/hal+json'

# The response:
{
  "svm": {
    "uuid": "03157e81-24c5-11e9-9ec1-005056bba643",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/03157e81-24c5-11e9-9ec1-005056bba643"
      }
    }
  }
},

```

```

"lun": {
  "uuid": "a60d9862-9bee-49a6-8162-20d2421bb1a6",
  "name": "/vol/vol1/lun1",
  "node": {
    "uuid": "7d8607ea-24c1-11e9-9ec1-005056bba643",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/7d8607ea-24c1-11e9-9ec1-005056bba643"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/luns/a60d9862-9bee-49a6-8162-20d2421bb1a6"
    }
  }
},
"igroup": {
  "uuid": "40d98b2c-24c5-11e9-9ec1-005056bba643",
  "name": "ig1",
  "os_type": "linux",
  "protocol": "mixed",
  "_links": {
    "self": {
      "href": "/api/protocols/san/igroups/40d98b2c-24c5-11e9-9ec1-005056bba643"
    }
  }
},
"logical_unit_number": 0,
"_links": {
  "self": {
    "href": "/api/protocols/san/lun-maps/a60d9862-9bee-49a6-8162-20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643"
  }
}
}

'''

=== Deleting a LUN map

----

```

```

# The API:
DELETE /api/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/protocols/san/lun-maps/a60d9862-
9bee-49a6-8162-20d2421bb1a6/40d98b2c-24c5-11e9-9ec1-005056bba643' -H
'accept: application/hal+json'
----

[[IDe517f5c30bbb24e4b08bf881edfc1a78]]
= Retrieve LUN maps

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/san/lun-maps`#

*Introduced In:* 9.6

Retrieves LUN maps.

== Related ONTAP commands

* `lun mapping show`
* xref:{relative_path}protocols_san_lun-maps_endpoint_overview.html[DOC
/protocols/san/lun-maps]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|lun.node.name
|string
|query
|False
a|Filter by lun.node.name

```

```
|lun.node.uuid  
|string  
|query  
|False  
a|Filter by lun.node.uuid
```

```
|lun.name  
|string  
|query  
|False  
a|Filter by lun.name
```

```
|lun.uuid  
|string  
|query  
|False  
a|Filter by lun.uuid
```

```
|svm.uuid  
|string  
|query  
|False  
a|Filter by svm.uuid
```

```
|svm.name  
|string  
|query  
|False  
a|Filter by svm.name
```

```
|igroup.protocol  
|string  
|query  
|False  
a|Filter by igroup.protocol
```

```
|igroup.name  
|string  
|query  
|False
```



```

a|Filter by igroup.name

|igroup.uuid
|string
|query
|False
a|Filter by igroup.uuid

|igroup.os_type
|string
|query
|False
a|Filter by igroup.os_type

|igroup.initiators
|string
|query
|False
a|Filter by igroup.initiators

|logical_unit_number
|integer
|query
|False
a|Filter by logical_unit_number

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean

```

```

|query
|False
a|The default is true for GET calls.  When set to false, only the number
of records is returned.

* Default value: 1


|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0


|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records

```

```

|integer
a|Number of records.

|records
|array[link:#lun_map[lun_map]]
a|

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "initiators": [
          "iqn.1998-01.com.corp.iscsi:name1"
        ],
        "name": "igroup1",
        "os_type": "string",
        "protocol": "string",
        "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
      },
      "logical_unit_number": "1",

```

```

    "lun": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "/vol/volume1/qtree1/lun1",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cf8aa42-8cd1-12e0-a11c-423468563412"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  ]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

```

```

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

```

```

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#igroup]
[.api-collapsible-fifth-title]
igroup

```

The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|initiators
|array[string]
a|The initiators that are members of the initiator group.

|name
|string
a|The name of the initiator group. Valid in POST.

|os_type
|string
a|The host operating system of the initiator group. All initiators in the
group should be hosts of the same operating system.

|protocol
|string
a|The protocols supported by the initiator group. This restricts the type
of initiators that can be added to the initiator group.

|uuid
|string
a|The unique identifier of the initiator group. Valid in POST.

|===

[#node]
[.api-collapsible-fifth-title]
node

The LUN node.

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name the LUN's node.

```

```

|uuid
|string
a|The unique identifier of the LUN node.

```

```

|===

```

```

[#lun]
[.api-collapsible-fifth-title]
lun

```

The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both.

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The fully qualified path name of the LUN composed of a "/vol" prefix,
the volume name, the (optional) qtree name, and file name of the LUN.
Valid in POST.

```



```
|node
|link:#node[node]
a|The LUN node.

|uuid
|string
a|The unique identifier of the LUN. Valid in POST.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#lun_map]
[.api-collapsible-fifth-title]
lun_map
```

A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator

group is many LUNs to many initiator groups.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|igroup
```

```
|link:#igroup[igroup]
```

a|The initiator group to which the LUN is mapped. Required in POST by supplying either the ``igroup.uuid``, ``igroup.name``, or both.

```
|logical_unit_number
```

```
|integer
```

a|The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value.

* example: 1

* Max value: 4095

* Min value: 0

* Introduced in: 9.6

* readCreate: 1

```
|lun
```

```
|link:#lun[lun]
```

a|The LUN to which the initiator group is mapped. Required in POST by supplying either the ``lun.uuid``, ``lun.name``, or both.

```
|svm
```

```
|link:#svm[svm]
```

```
a|
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID88275bf66e015e67d980d3e9ae2ed639]]
= Create a LUN map
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/protocols/san/lun-maps`#
```

Introduced In: 9.6

Creates a LUN map.

== Required properties

- * `svm.uuid` or `svm.name` - Existing SVM in which to create the LUN map.
- * `igroup.uuid` or `igroup.name` - Existing initiator group to map to the specified LUN.
- * `lun.uuid` or `lun.name` - Existing LUN to map to the specified initiator group.

== Default property values

If not specified in POST, the following default property values are assigned.

- * `logical_unit_number` - If no value is provided, ONTAP assigns the lowest available value.

== Related ONTAP commands

- * `lun mapping create`

== Learn more

- * xref:{relative_path}protocols_san_lun-maps_endpoint_overview.html[DOC

```
/protocols/san/lun-maps]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|return_records
```

```
|boolean
```

```
|query
```

```
|False
```

```
a|The default is false. If set to true, the records are returned.
```

```
* Default value:
```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|igroup
```

```
|link:#igroup[igroup]
```

```
a|The initiator group to which the LUN is mapped. Required in POST by  
supplying either the `igroup.uuid`, `igroup.name`, or both.
```

```
|logical_unit_number
```

```
|integer
```

```
a|The logical unit number assigned to the LUN when mapped to the specified  
initiator group. The number is used to identify the LUN to initiators in  
the initiator group when communicating through Fibre Channel Protocol or  
iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest  
available value.
```

```
* example: 1
* Max value: 4095
* Min value: 0
* Introduced in: 9.6
* readCreate: 1
```

```
|lun
|link:#lun[lun]
a|The LUN to which the initiator group is mapped. Required in POST by
supplying either the `lun.uuid`, `lun.name`, or both.
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

```
.Example request
```

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{
  "igroup": {
    "initiators": [
      "iqn.1998-01.com.corp.iscsi:name1"
    ],
    "name": "igroup1",
    "os_type": "string",
    "protocol": "string",
    "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
  },
  "logical_unit_number": "1",
  "lun": {
    "name": "/vol/volume1/qtrees1/lun1",
    "node": {
      "name": "node1",
      "uuid": "1cf8aa42-8cd1-12e0-a11c-423468563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
```

```
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====
```

== Response

Status: 201, Created

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|num_records
|integer
a|Number of records.

|records
|array[link:#lun_map[lun_map]]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "records": [
    {
      "igroup": {
        "initiators": [
          "iqn.1998-01.com.corp.iscsi:name1"
        ],
        "name": "igroup1",
        "os_type": "string",
        "protocol": "string",
        "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
      },
      "logical_unit_number": "1",
      "lun": {
        "name": "/vol/volume1/qtreet1/lun1",
```

```

    "node": {
      "name": "node1",
      "uuid": "1cf8aa42-8cd1-12e0-a11c-423468563412"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
|
| 1254207
| The LUN is already mapped to the same initiator group.
|
| 2621462
| The specific SVM does not exist.
|
| 2621706
| Both the SVM UUID and SVM name were supplied, but don't refer to the
same SVM.
|
| 2621707
| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.
|
| 5374053
| The LUN is the destination of an ongoing restore operation and is
inaccessible for I/O and management. Wait for the restore to complete and
try the command again.
|
| 5374238
| The operation is not allowed on a LUN in a Snapshot copy.
|
| 5374316

```



```

| A LUN move operation is in progress on the source LUN.

| 5374329
| A LUN of class `vvol` cannot be mapped.

| 5374573
| A node has no interface configured with the iSCSI or Fibre Channel
protocols for the specified SVM.

| 5374574
| Multiple nodes have no interface configured with the iSCSI or Fibre
Channel protocols for the specified SVM.

| 5374581
| A node has no interface configured with the iSCSI protocol for the
specified SVM.

| 5374582
| Multiple nodes have no interface configured with the iSCSI protocol for
the specified SVM.

| 5374583
| A node has no interface configured with the Fibre Channel protocol for
the specified SVM.

| 5374584
| Multiple nodes have no interface configured with the Fibre Channel
protocol for the specified SVM.

| 5374901
| Either `lun.uuid` or `lun.name` must be provided to create a LUN map.

| 5374902
| Either `igroup.uuid` or `igroup.name` must be provided to create a LUN
map.
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

[#href]
[.api-collapsible-fifth-title]

```

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

|string

a|

|===

[#_links]

[.api-collapsible-fifth-title]

_links

[#igroup]

[.api-collapsible-fifth-title]

igroup

The initiator group to which the LUN is mapped. Required in POST by supplying either the `igroup.uuid`, `igroup.name`, or both.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|initiators

|array[string]

a|The initiators that are members of the initiator group.

|name

|string

a|The name of the initiator group. Valid in POST.

|os_type

|string

a|The host operating system of the initiator group. All initiators in the group should be hosts of the same operating system.

```
|protocol
|string
a|The protocols supported by the initiator group. This restricts the type
of initiators that can be added to the initiator group.
```

```
|uuid
|string
a|The unique identifier of the initiator group. Valid in POST.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

The LUN node.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The name the LUN's node.
```

```
|uuid
|string
a|The unique identifier of the LUN node.
```

```
|===
```

```
[#lun]
[.api-collapsible-fifth-title]
lun
```

The LUN to which the initiator group is mapped. Required in POST by supplying either the ``lun.uuid``, ``lun.name``, or both.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The fully qualified path name of the LUN composed of a "/vol" prefix,
the volume name, the (optional) qtree name, and file name of the LUN.
Valid in POST.

|node
|link:#node[node]
a|The LUN node.

|uuid
|string
a|The unique identifier of the LUN. Valid in POST.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

```

|===

```
[#lun_map]
[.api-collapsible-fifth-title]
lun_map
```

A LUN map is an association between a LUN and an initiator group. When a LUN is mapped to an initiator group, the initiator group's initiators are granted access to the LUN. The relationship between a LUN and an initiator group is many LUNs to many initiator groups.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|igroup
```

```
|link:#igroup[igroup]
```

a|The initiator group to which the LUN is mapped. Required in POST by supplying either the ``igroup.uuid``, ``igroup.name``, or both.

```
|logical_unit_number
```

```
|integer
```

a|The logical unit number assigned to the LUN when mapped to the specified initiator group. The number is used to identify the LUN to initiators in the initiator group when communicating through Fibre Channel Protocol or iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest available value.

- * example: 1
- * Max value: 4095
- * Min value: 0
- * Introduced in: 9.6
- * readCreate: 1

```
|lun
```

```
|link:#lun[lun]
```

a|The LUN to which the initiator group is mapped. Required in POST by supplying either the ``lun.uuid``, ``lun.name``, or both.

```

|svm
|link:#svm[svm]
a|

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string

```

```

a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID2e211b4b8385f2b44c2ac7e6f0062655]]
= Delete a LUN map

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}`#

*Introduced In:* 9.6

Deletes a LUN map.

== Related ONTAP commands

* `lun mapping delete`

== Learn more

* xref:{relative_path}protocols_san_lun-maps_endpoint_overview.html[DOC
/protocols/san/lun-maps]

== Parameters

[cols=5*,options=header]
|===

|Name

```

```

|Type
|In
|Required
|Description

|lun.uuid
|string
|path
|True
a|The unique identifier of the LUN.

|igroup.uuid
|string
|path
|True
a|The unique identifier of the igroup.

|===

== Response

```

Status: 200, Ok

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 5374875
| The LUN was not found.
|===

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```



```

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```

[[ID6cf7746019d7d84973586eb2a041d687]]
= Retrieve a LUN map

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/protocols/san/lun-maps/{lun.uuid}/{igroup.uuid}`#

*Introduced In:* 9.6

Retrieves a LUN map.

== Related ONTAP commands

* `lun mapping show`

== Learn more

* xref:{relative_path}protocols_san_lun-maps_endpoint_overview.html[DOC
/protocols/san/lun-maps]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|lun.uuid
|string
|path
|True
a|The unique identifier of the LUN.

|igroup.uuid
|string
|path
|True
a|The unique identifier of the igroup.

|fields
|array[string]

```

```
|query
|False
a|Specify the fields to return.

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|igroup
|link:#igroup[igroup]
a|The initiator group to which the LUN is mapped. Required in POST by
supplying either the `igroup.uuid`, `igroup.name`, or both.

|logical_unit_number
|integer
a|The logical unit number assigned to the LUN when mapped to the specified
initiator group. The number is used to identify the LUN to initiators in
the initiator group when communicating through Fibre Channel Protocol or
iSCSI. Optional in POST; if no value is provided, ONTAP assigns the lowest
available value.

* example: 1
* Max value: 4095
* Min value: 0
* Introduced in: 9.6
* readCreate: 1

|lun
|link:#lun[lun]
a|The LUN to which the initiator group is mapped. Required in POST by
supplying either the `lun.uuid`, `lun.name`, or both.
```

```
|svm
|link:#svm[svm]
a|
```

```
|===
```

.Example response

[%collapsible%closed]

```
=====
```

[source,json,subs==+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "igroup": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "initiators": [
      "iqn.1998-01.com.corp.iscsi:namel"
    ],
    "name": "igroup1",
    "os_type": "string",
    "protocol": "string",
    "uuid": "1ad8544d-8cd1-91e0-9e1c-723478563412"
  },
  "logical_unit_number": "1",
  "lun": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "/vol/volume1/qtree1/lun1",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
}
```

```

    "name": "node1",
    "uuid": "1cf8aa42-8cd1-12e0-a11c-423468563412"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description

| 5374852
| The initiator group was not found.

| 5374875
| The LUN was not found.
|===

```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]

```

```
[.api-collapsible-fifth-title]
```

```
_links
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|self
```

```
|link:#href[href]
```

```
a|
```

```
|===
```

```
[#igroup]
```

```
[.api-collapsible-fifth-title]
```

```
igroup
```

The initiator group to which the LUN is mapped. Required in POST by supplying either the ``igroup.uuid``, ``igroup.name``, or both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|initiators
```

```
|array[string]
```

a|The initiators that are members of the initiator group.

```
|name
```

```
|string
```

a|The name of the initiator group. Valid in POST.

```
|os_type
```

```
|string
```

a|The host operating system of the initiator group. All initiators in the

group should be hosts of the same operating system.

|protocol

|string

a|The protocols supported by the initiator group. This restricts the type of initiators that can be added to the initiator group.

|uuid

|string

a|The unique identifier of the initiator group. Valid in POST.

|===

[#node]

[.api-collapsible-fifth-title]

node

The LUN node.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name the LUN's node.

|uuid

|string

a|The unique identifier of the LUN node.

|===

```
[#lun]
[.api-collapsible-fifth-title]
lun
```

The LUN to which the initiator group is mapped. Required in POST by supplying either the `lun.uuid`, `lun.name`, or both.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The fully qualified path name of the LUN composed of a "/vol" prefix,
the volume name, the (optional) qtree name, and file name of the LUN.
Valid in POST.
```

```
|node
|link:#node[node]
a|The LUN node.
```

```
|uuid
|string
a|The unique identifier of the LUN. Valid in POST.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage LUNs

:leveloffset: +1

[[IDf72d1b9e16c0a62a4511d305a6eb8456]]
= Storage luns endpoint overview


== Overview

```

A LUN is the logical representation of storage in a storage area network (SAN).

The LUN REST API allows you to create, update, delete, and discover LUNs.

In ONTAP, a LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, and moved to a different volume. LUNs support the assignment of a quality of service (QoS) policy for performance management or a QoS policy can be assigned to the volume containing the LUN. See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the FC Protocol or a TCP/IP network using iSCSI.

== Performance monitoring

Performance of a LUN can be monitored by observing the ``metric.++`` and ``statistics.++`` properties. These properties show the performance of a LUN in terms of IOPS, latency and throughput. The ``metric.++`` properties denote an average whereas ``statistics.++`` properties denote a real-time monotonically increasing value aggregated across all nodes.

== Examples

=== Creating a LUN

This example creates a 300 gigabyte, thin-provisioned LUN in SVM `_svm1_`, volume `_voll_`, configured for use by `_linux_` initiators. The ``return_records`` query parameter is used to retrieve properties of the newly created LUN in the POST response.

The API:

POST /api/storage/luns

The call:

```
curl -X POST 'https://<mgmt-ip>/api/storage/luns?return_records=true' -H
'accept: application/hal+json' -d '{ "svm": { "name": "svm1" }, "os_type":
"linux", "space": { "size": "300G" }, "name" : "/vol/voll/lun1" }'
```

```
# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "5a24ae5b-28af-47fb-b129-5adf6cfba0a6",
      "svm": {
        "uuid": "6bf967fd-2a1c-11e9-b682-005056bbc17d",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/6bf967fd-2a1c-11e9-b682-005056bbc17d"
          }
        }
      },
      "name": "/vol/vol1/lun1",
      "location": {
        "logical_unit": "lun1",
        "volume": {
          "uuid": "71cd0dba-2a1c-11e9-b682-005056bbc17d",
          "name": "vol1",
          "_links": {
            "self": {
              "href": "/api/storage/volumes/71cd0dba-2a1c-11e9-b682-005056bbc17d"
            }
          }
        }
      },
      "class": "regular",
      "enabled": true,
      "os_type": "linux",
      "serial_number": "wf0Iq+N4uck3",
      "space": {
        "size": 322163441664,
        "used": 0,
        "guarantee": {
          "requested": false,
          "reserved": false
        }
      },
      "status": {
        "container_state": "online",
        "read_only": false,
        "state": "online"
      }
    },
  ],
}
```

```

    "_links": {
      "self": {
        "href": "/api/storage/luns/5a24ae5b-28af-47fb-b129-5adf6cfba0a6"
      }
    }
  }
]
}
-----
'''

```

=== Updating a LUN

This example sets the `comment` property of a LUN.

```

-----

# The API:
PATCH /api/storage/luns/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/luns/5a24ae5b-28af-47fb-b129-5adf6cfba0a6' -H 'accept: application/hal+json' -d '{ "comment": "Data for the finance department." }'
-----

'''

```

=== Retrieving LUNs

This example retrieves summary information for all online LUNs in SVM `_svm1_`. The `svm.name` and `status.state` query parameters are used to find the desired LUNs.

```

-----

# The API:
GET /api/storage/luns

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/luns?svm.name=svm1&status.state=online' -H 'accept: application/hal+json'

# The response:
{

```

```
"records": [  
  {  
    "uuid": "5a24ae5b-28af-47fb-b129-5adf6cfba0a6",  
    "svm": {  
      "name": "svm1"  
    },  
    "name": "/vol/vol1/lun1",  
    "status": {  
      "state": "online"  
    },  
    "_links": {  
      "self": {  
        "href": "/api/storage/luns/5a24ae5b-28af-47fb-b129-5adf6cfba0a6"  
      }  
    }  
  },  
  {  
    "uuid": "c903a978-9bac-4ce9-8237-4a3ba8b13f08",  
    "svm": {  
      "name": "svm1"  
    },  
    "name": "/vol/vol1/lun2",  
    "status": {  
      "state": "online"  
    },  
    "_links": {  
      "self": {  
        "href": "/api/storage/luns/c903a978-9bac-4ce9-8237-4a3ba8b13f08"  
      }  
    }  
  },  
  {  
    "uuid": "7faf0a9e-0a47-4876-8318-3638d5da16bf",  
    "svm": {  
      "name": "svm1"  
    },  
    "name": "/vol/vol2/lun3",  
    "status": {  
      "state": "online"  
    },  
    "_links": {  
      "self": {  
        "href": "/api/storage/luns/7faf0a9e-0a47-4876-8318-3638d5da16bf"  
      }  
    }  
  }  
]
```



```

],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/luns?svm.name=svm1&status.state=online"
  }
}
}
}

```

'''

=== Retrieving details for a specific LUN

In this example, the ``fields`` query parameter is used to request all fields, including advanced fields, that would not otherwise be returned by default for the LUN.

The API:

GET /api/storage/luns/{uuid}

The call:

curl -X GET 'https://<mgmt-ip>/api/storage/luns/5a24ae5b-28af-47fb-b129-5adf6cfba0a6?fields=**' -H 'accept: application/hal+json'

The response:

```

{
  "uuid": "5a24ae5b-28af-47fb-b129-5adf6cfba0a6",
  "svm": {
    "uuid": "6bf967fd-2a1c-11e9-b682-005056bbc17d",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/6bf967fd-2a1c-11e9-b682-005056bbc17d"
      }
    }
  },
  "name": "/vol/vol1/lun1",
  "location": {
    "logical_unit": "lun1",
    "volume": {
      "uuid": "71cd0dba-2a1c-11e9-b682-005056bbc17d",
      "name": "vol1",
      "_links": {

```

```

    "self": {
      "href": "/api/storage/volumes/71cd0dba-2a1c-11e9-b682-005056bbc17d"
    }
  },
  "auto_delete": false,
  "class": "regular",
  "comment": "Data for the finance department.",
  "enabled": true,
  "lun_maps": [
    {
      "logical_unit_number": 0,
      "igroup": {
        "uuid": "2b9d57e1-2a66-11e9-b682-005056bbc17d",
        "name": "ig1",
        "_links": {
          "self": {
            "href": "/api/protocols/san/igroups/2b9d57e1-2a66-11e9-b682-005056bbc17d"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/protocols/san/lun-maps/5a24ae5b-28af-47fb-b129-5adf6cfba0a6/2b9d57e1-2a66-11e9-b682-005056bbc17d"
        }
      }
    }
  ],
  "os_type": "linux",
  "serial_number": "wf0Iq+N4uck3",
  "space": {
    "size": 322163441664,
    "used": 0,
    "guarantee": {
      "requested": false,
      "reserved": false
    }
  },
  "metric": {
    "timestamp": "2019-04-09T05:50:15Z",
    "duration": "PT15S",
    "status": "ok",

```

```
"latency": {
  "other": 0,
  "total": 0,
  "read": 0,
  "write": 0
},
"iops": {
  "read": 0,
  "write": 0,
  "other": 0,
  "total": 0
},
"throughput": {
  "read": 0,
  "write": 0,
  "other": 0,
  "total": 0
}
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "latency_raw": {
    "other": 38298,
    "total": 38298,
    "read": 0,
    "write": 0
  },
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,
    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"status": {
  "container_state": "online",
  "mapped": true,
  "read_only": false,
  "state": "online"
}
```

```

},
"_links": {
  "self": {
    "href": "/api/storage/luns/5a24ae5b-28af-47fb-b129-5adf6cfba0a6?fields=**"
  }
}
}
}
----

'''

```

== Cloning LUNs

A clone of a LUN is an independent "copy" of the LUN that shares unchanged data blocks with the original. As blocks of the source and clone are modified, unique blocks are written for each. LUN clones can be created quickly and consume very little space initially. They can be created for the purpose of back-up, or to replicate data for multiple consumers.

Space reservations can be set for the LUN clone independent of the source LUN by setting the ``space.guarantee.requested`` property in a POST or PATCH request.

A LUN clone can also be set to auto-delete by setting the ``auto_delete`` property. If the LUN's volume is configured for automatic deletion, LUNs that have auto-delete enabled are deleted when a volume is nearly full to reclaim a target amount of free space in the volume.

== Examples

=== Creating a new LUN clone

You create a new LUN clone as you create any LUN - a POST request to `xref:{relative_path}getting_started_with_the_ontap_rest_api.html#/SAN/lun_create[/storage/luns]`. Set ``clone.source.uuid`` or ``clone.source.name`` to identify the source LUN from which the clone is created. The LUN clone and its source must reside in the same volume.

The source LUN can reside in a Snapshot copy, in which case the ``clone.source.name`` field must be used to identify it. Add ``/.snapshot/<snapshot_name>`` to the path after the volume name to identify the Snapshot copy. For example ``/vol/vol1/.snapshot/snap1/lun1``.

By default, new LUN clones do not inherit the QoS policy of the source LUN; a QoS policy should be set for the clone by setting the ``qos_policy`` property.

The API:

POST /api/storage/luns

The call:

```
curl -X POST 'https://<mgmt-ip>/api/storage/luns' -H 'accept:
application/hal+json' -d '{ "svm": { "name": "svm1" }, "name":
"/vol/vol1/lun2clone1", "clone": { "source": { "name": "/vol/vol1/lun2" }
}, "qos_policy": { "name": "qos1" } }'
```

=== Over-writing an existing LUN's data as a clone of another

You can overwrite an existing LUN as a clone of another, using a PATCH request to

xref:{relative_path}getting_started_with_the_ontap_rest_api.html#/SAN/lun_modify[/storage/luns/{uuid}]. Set the ``clone.source.uuid`` or ``clone.source.name`` property to identify the source LUN from which the clone data is taken. The LUN clone and its source must reside in the same volume.

When used in a PATCH request, the patched LUN's data is overwritten as a clone of the source. The following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: ``class``, ``auto_delete``, ``lun_maps``, ``serial_number``, ``status.state``, and ``uuid``.

Persistent reservations for the updated LUN are also preserved.

The API:

PATCH /api/storage/luns/{uuid}

The call:

```
curl -X PATCH 'https://<mgmt-ip>/api/storage/luns/5a24ae5b-28af-47fb-b129-
5adf6cfba0a6' -H 'accept: application/hal+json' -d '{ "clone": { "source":
{ "name": "/vol/vol1/lun2" } } }'
```

...

== Moving LUNs between volumes

You move a LUN between volumes by using a PATCH request to

xref:{relative_path}getting_started_with_the_ontap_rest_api.html#/SAN/lun_

modify[/storage/luns/{uuid}]. Set the volume portion of the fully qualified LUN path `name` property, `path.volume.uuid`, or `path.volume.name` property to a different volume than the LUN's current volume. Moving a LUN between volumes is an asynchronous activity. A successful request returns a response of 200 synchronously, which indicates that the movement has been successfully queued. The LUN object can then be further polled with a GET request to <<lun_get,`+/storage/luns/{uuid}+`>> to monitor the status of the movement.

The `movement` sub-object of the LUN object is populated while a LUN movement is in progress and for two minutes following completion of a movement.

== Examples

=== Starting a LUN movement

The API:

PATCH /api/storage/luns/{uuid}

The call:

```
curl -X PATCH 'https://<mgmt-ip>/api/storage/luns/7faf0a9e-0a47-4876-8318-3638d5da16bf' -H 'accept: application/hal+json' -d '{ "name": "/vol/vol1/lun3" }'
```

=== Checking on the status of the LUN movement

The API:

GET /api/storage/luns/{uuid}

The call:

```
curl -X GET 'https://<mgmt-ip>/api/storage/luns/7faf0a9e-0a47-4876-8318-3638d5da16bf?fields=movement' -H 'accept: application/hal+json'
```

The response:

```
{
  "uuid": "7faf0a9e-0a47-4876-8318-3638d5da16bf",
  "name": "/vol/vol1/lun3",
  "movement": {
    "paths": {
      "destination": "/vol/vol1/lun3",
```

```

    "source": "/vol/vol2/lun3"
  },
  "progress": {
    "elapsed": 1,
    "percent_complete": 0,
    "state": "preparing",
    "volume_snapshot_blocked": false
  }
},
"_links": {
  "self": {
    "href": "/api/storage/luns/7faf0a9e-0a47-4876-8318-3638d5da16bf"
  }
}
}
}
----

'''

=== Deleting a LUN

----

# The API:
DELETE /api/storage/luns/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/storage/luns/c903a978-9bac-4ce9-8237-4a3ba8b13f08' -H 'accept: application/hal+json'
----

'''

[[ID8a3107483485c625e0da2487fdf904fb]]
= Retrieve LUNs

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/storage/luns`#

*Introduced In:* 9.6

Retrieves LUNs.

```

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

- * `auto_delete`
- * `lun_maps.+++`
- * `movement.+++`
- * `status.mapped`
- * `statistics.+++`
- * `metric.+++`

== Related ONTAP commands

- * `lun mapping show`
- * `lun move show`
- * `lun show`
- * `volume file clone show-autodelete`

== Learn more

- * [xref:{relative_path}storage_luns_endpoint_overview.html\[DOC /storage/luns\]](#)

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

create_time
string
query
False
a Filter by create_time

- * Introduced in: 9.7


```
|location.logical_unit
|string
|query
|False
a|Filter by location.logical_unit
```

```
|location.qtree.name
|string
|query
|False
a|Filter by location.qtree.name
```

```
|location.qtree.id
|integer
|query
|False
a|Filter by location.qtree.id
```

```
|location.volume.uuid
|string
|query
|False
a|Filter by location.volume.uuid
```

```
|location.volume.name
|string
|query
|False
a|Filter by location.volume.name
```

```
|enabled
|boolean
|query
|False
a|Filter by enabled
```

```
|qos_policy.uuid
|string
|query
|False
```

```
a|Filter by qos_policy.uuid

|qos_policy.name
|string
|query
|False
a|Filter by qos_policy.name

|svm.uuid
|string
|query
|False
a|Filter by svm.uuid

|svm.name
|string
|query
|False
a|Filter by svm.name

|movement.max_throughput
|string
|query
|False
a|Filter by movement.max_throughput

|movement.paths.destination
|string
|query
|False
a|Filter by movement.paths.destination

|movement.paths.source
|string
|query
|False
a|Filter by movement.paths.source

|movement.progress.volume_snapshot_blocked
|boolean
```

```
|query
|False
a|Filter by movement.progress.volume_snapshot_blocked
```

```
|movement.progress.state
|string
|query
|False
a|Filter by movement.progress.state
```

```
|movement.progress.elapsed
|integer
|query
|False
a|Filter by movement.progress.elapsed
```

```
|movement.progress.failure.code
|string
|query
|False
a|Filter by movement.progress.failure.code
```

```
|movement.progress.failure.message
|string
|query
|False
a|Filter by movement.progress.failure.message
```

```
|movement.progress.percent_complete
|integer
|query
|False
a|Filter by movement.progress.percent_complete
```

```
|statistics.iops_raw.total
|integer
|query
|False
a|Filter by statistics.iops_raw.total
```

```
* Introduced in: 9.7
```

```
|statistics.iops_raw.read
|integer
|query
|False
a|Filter by statistics.iops_raw.read
```

* Introduced in: 9.7

```
|statistics.iops_raw.other
|integer
|query
|False
a|Filter by statistics.iops_raw.other
```

* Introduced in: 9.7

```
|statistics.iops_raw.write
|integer
|query
|False
a|Filter by statistics.iops_raw.write
```

* Introduced in: 9.7

```
|statistics.latency_raw.total
|integer
|query
|False
a|Filter by statistics.latency_raw.total
```

* Introduced in: 9.7

```
|statistics.latency_raw.read
|integer
|query
|False
a|Filter by statistics.latency_raw.read
```

* Introduced in: 9.7

```
|statistics.latency_raw.other  
|integer  
|query  
|False  
a|Filter by statistics.latency_raw.other
```

* Introduced in: 9.7

```
|statistics.latency_raw.write  
|integer  
|query  
|False  
a|Filter by statistics.latency_raw.write
```

* Introduced in: 9.7

```
|statistics.timestamp  
|string  
|query  
|False  
a|Filter by statistics.timestamp
```

* Introduced in: 9.7

```
|statistics.status  
|string  
|query  
|False  
a|Filter by statistics.status
```

* Introduced in: 9.7

```
|statistics.throughput_raw.total  
|integer  
|query  
|False  
a|Filter by statistics.throughput_raw.total
```

* Introduced in: 9.7

```
|statistics.throughput_raw.read  
|integer
```

```
|query
|False
a|Filter by statistics.throughput_raw.read

* Introduced in: 9.7

|statistics.throughput_raw.other
|integer
|query
|False
a|Filter by statistics.throughput_raw.other

* Introduced in: 9.7

|statistics.throughput_raw.write
|integer
|query
|False
a|Filter by statistics.throughput_raw.write

* Introduced in: 9.7

|serial_number
|string
|query
|False
a|Filter by serial_number

|auto_delete
|boolean
|query
|False
a|Filter by auto_delete

|uuid
|string
|query
|False
a|Filter by uuid

|os_type
```

```
|string
|query
|False
a|Filter by os_type
```

```
|metric.timestamp
|string
|query
|False
a|Filter by metric.timestamp
```

* Introduced in: 9.7

```
|metric.duration
|string
|query
|False
a|Filter by metric.duration
```

* Introduced in: 9.7

```
|metric.throughput.total
|integer
|query
|False
a|Filter by metric.throughput.total
```

* Introduced in: 9.7

```
|metric.throughput.read
|integer
|query
|False
a|Filter by metric.throughput.read
```

* Introduced in: 9.7

```
|metric.throughput.other
|integer
|query
|False
a|Filter by metric.throughput.other
```

* Introduced in: 9.7

```
|metric.throughput.write  
|integer  
|query  
|False  
a|Filter by metric.throughput.write
```

* Introduced in: 9.7

```
|metric.latency.total  
|integer  
|query  
|False  
a|Filter by metric.latency.total
```

* Introduced in: 9.7

```
|metric.latency.read  
|integer  
|query  
|False  
a|Filter by metric.latency.read
```

* Introduced in: 9.7

```
|metric.latency.other  
|integer  
|query  
|False  
a|Filter by metric.latency.other
```

* Introduced in: 9.7

```
|metric.latency.write  
|integer  
|query  
|False  
a|Filter by metric.latency.write
```

* Introduced in: 9.7


```
|metric.status
|string
|query
|False
a|Filter by metric.status
```

* Introduced in: 9.7

```
|metric.iops.total
|integer
|query
|False
a|Filter by metric.iops.total
```

* Introduced in: 9.7

```
|metric.iops.read
|integer
|query
|False
a|Filter by metric.iops.read
```

* Introduced in: 9.7

```
|metric.iops.other
|integer
|query
|False
a|Filter by metric.iops.other
```

* Introduced in: 9.7

```
|metric.iops.write
|integer
|query
|False
a|Filter by metric.iops.write
```

* Introduced in: 9.7

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|space.size
|integer
|query
|False
a|Filter by space.size
```

```
|space.used
|integer
|query
|False
a|Filter by space.used
```

```
|space.guarantee.reserved
|boolean
|query
|False
a|Filter by space.guarantee.reserved
```

```
|space.guarantee.requested
|boolean
|query
|False
a|Filter by space.guarantee.requested
```

```
|class
|string
|query
|False
a|Filter by class
```

```
|status.read_only
|boolean
|query
|False
a|Filter by status.read_only
```

```
|status.mapped
|boolean
|query
|False
a|Filter by status.mapped
```

```
|status.container_state
|string
|query
|False
a|Filter by status.container_state
```

```
|status.state
|string
|query
|False
a|Filter by status.state
```

```
|lun_maps.logical_unit_number
|integer
|query
|False
a|Filter by lun_maps.logical_unit_number
```

```
|lun_maps.igroup.name
|string
|query
|False
a|Filter by lun_maps.igroup.name
```

```
|lun_maps.igroup.uuid
|string
|query
```

```

|False
a|Filter by lun_maps.igroup.uuid

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.

```

Default direction is 'asc' for ascending.

|===

== Response

Status: 200, Ok

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|num_records

|integer

a|Number of records.

|records

|array[link:#lun[lun]]

a|

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "_links": {

 "next": {

 "href": "/api/resourcelink"

 },

 "self": {

 "href": "/api/resourcelink"

 }

 },

 "records": [

 {

 "_links": {

```

    "self": {
      "href": "/api/resourcelink"
    },
    "class": "string",
    "comment": "string",
    "create_time": "2018-06-04T19:00:00Z",
    "location": {
      "logical_unit": "lun1",
      "qtree": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": "1",
        "name": "qt1"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    },
    "lun_maps": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "igroup": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        },
        "logical_unit_number": 0
      }
    ]
  }

```

```

],
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"movement": {
  "max_throughput": "string",
  "paths": {
    "destination": "/vol/vol1/lun1",
    "source": "/vol/vol2/lun2"
  },
  "progress": {
    "elapsed": 0,
    "failure": {
      "code": "4",
      "message": "Destination volume is offline."
    },
    "percent_complete": 0,
    "state": "string"
  }
},
"name": "/vol/volume1/qtree1/lun1",
"os_type": "string",
"qos_policy": {
  "_links": {

```

```

    "self": {
      "href": "/api/resourcelink"
    },
    },
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "serial_number": "string",
  "space": {
    "size": "1073741824",
    "used": 0
  },
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

```



```

]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block

```

```

====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type

```

|Description

|self
|link:#href[href]
a|

|===

[#source]
[.api-collapsible-fifth-title]
source

The source LUN for a LUN clone operation. This can be specified using property ``clone.source.uuid`` or ``clone.source.name``. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

[#clone]
[.api-collapsible-fifth-title]
clone

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: ``auto_delete``, ``qos_policy``, and ``space.guarantee.requested``.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: ``class``, ``auto_delete``, ``lun_maps``, ``serial_number``, ``status.state``, and ``uuid``.

Persistent reservations for the patched LUN are also preserved.

[#qtree]
[.api-collapsible-fifth-title]
qtree

The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties ``name`` and ``location.qtree.name`` and/or ``location.qtree.uuid`` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|id
```

```
|integer
```

```
a|The identifier for the qtree, unique within the qtree's volume.
```

```
|name
```

```
|string
```

```
a|The name of the qtree.
```

```
|===
```

```
[#volume]
```

```
[.api-collapsible-fifth-title]
```

```
volume
```

The volume in which the LUN is located. Valid in POST and PATCH.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume.
```

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|logical_unit
|string
a|The base name component of the LUN. Valid in POST and PATCH.
```

If properties `name` and `location.logical_unit` are specified in the same request, they must refer to the base name.

A PATCH that modifies the base name of the LUN is considered a rename operation.

```
|qtree
|link:#qtree[qtree]
a|The qtree in which the LUN is optionally located. Valid in POST and
PATCH.
```

If properties `name` and `location.qtree.name` and/or
`location.qtree.uuid` are specified in the same request, they must refer
to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename
operation.

```
|volume
|link:#volume[volume]
a|The volume in which the LUN is located. Valid in POST and PATCH.
```

If properties `name` and `location.volume.name` and/or
`location.volume.uuid` are specified in the same request, they must refer
to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN
movement operation.

```
|===
```

```
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

The initiator group to which the LUN is mapped.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
```

a|The name of the initiator group.

|uuid

|string

a|The unique identifier of the initiator group.

|===

[#lun_maps]

[.api-collapsible-fifth-title]

lun_maps

A LUN map with which the LUN is associated.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|igroup

|link:#igroup[igroup]

a|The initiator group to which the LUN is mapped.

|logical_unit_number

|integer

a|The logical unit number assigned to the LUN for initiators in the initiator group.

|===

[#iops]

[.api-collapsible-fifth-title]

iops

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency]
```

```
[.api-collapsible-fifth-title]
```

latency

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```


a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#throughput]

[.api-collapsible-fifth-title]

throughput

The rate of throughput bytes per second observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

```

|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#metric]
[.api-collapsible-fifth-title]
metric

Performance numbers, such as IOPS latency and throughput.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status

```

```
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.
```

```
|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#paths]
[.api-collapsible-fifth-title]
paths
```

The fully qualified LUN path names involved in the LUN movement.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination
|string
a|The fully qualified path of the LUN movement destination composed of a
"/vol" prefix, the volume name, the (optional) qtree name, and base name
of the LUN.
```

```
|source
|string
a|The fully qualified path of the LUN movement source composed of a "/vol"
prefix, the volume name, the (optional) qtree name, and base name of the
LUN.
```

```
|===
```

```
[#failure]
[.api-collapsible-fifth-title]
failure
```

Error information provided if the asynchronous LUN movement operation fails.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|The error code.
```

```
|message
|string
a|The error message.
```

```
|===
```

```
[#progress]
[.api-collapsible-fifth-title]
progress
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

|elapsed
|integer
a|The amount of time, in seconds, that has elapsed since the start of the LUN movement.

|failure
|link:#failure[failure]
a|Error information provided if the asynchronous LUN movement operation fails.

|percent_complete
|integer
a|The percentage complete of the LUN movement.

|state
|string
a|The state of the LUN movement.

Valid in PATCH when an LUN movement is active. Set to `_paused_` to pause a LUN movement. Set to `_replicating_` to resume a paused LUN movement.

|volume_snapshot_blocked
|boolean
a|This property reports if volume Snapshot copies are blocked by the LUN movement. This property can be polled to identify when volume Snapshot copies can be resumed after beginning a LUN movement.

|===

[#movement]
[.api-collapsible-fifth-title]
movement

This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property ``name``, ``location.volume.uuid``, or ``location.volume.name``. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH

request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.

There is added cost to retrieving property values for `movement`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|max_throughput
```

```
|string
```

a|The maximum data throughput that should be utilized in support of the LUN movement. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.

If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.

* Introduced in: 9.6

```
|paths
```

```
|link:#paths[paths]
```

a|The fully qualified LUN path names involved in the LUN movement.

```
|progress
|link:#progress[progress]
a|
```

```
|===
```

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property ``qos_policy.uuid`` and ``qos_policy.name`` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property ``qos_policy.name`` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set this property to an empty string ("") in a PATCH request. Valid in POST and PATCH.
```

```
|uuid
|string
a|The unique identifier of the QoS policy. Valid in POST and PATCH.
```

```
|===
```

```
[#guarantee]
[.api-collapsible-fifth-title]
guarantee
```

Properties that request and report the space guarantee for the LUN.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|requested
|boolean
```

a|The requested space reservation policy for the LUN. If `_true_`, a space reservation is requested for the LUN; if `_false_`, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH.

```
|reserved
|boolean
```

a|Reports if the LUN is space guaranteed.

If `_true_`, a space guarantee is requested and the containing volume and aggregate support the request. If `_false_`, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request.

```
|===
```

```
[#space]
[.api-collapsible-fifth-title]
space
```

The storage space related properties of the LUN.

```
[cols=3*,options=header]
|===
|Name
|Type
```


|Description

|guarantee

|link:#guarantee[guarantee]

a|Properties that request and report the space guarantee for the LUN.

|size

|integer

a|The total provisioned size of the LUN. The LUN size can be increased but not be made smaller using the REST interface.

The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform and the available space in the containing volume and aggregate.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* example: 1073741824

* Max value: 140737488355328

* Min value: 4096

* Introduced in: 9.6

|used

|integer

a|The amount of space consumed by the main data stream of the LUN.

This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* readOnly: 1

* Introduced in: 9.6

|===

```
[#iops_raw]
[.api-collapsible-fifth-title]
iops_raw
```

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be  
metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

```
throughput_raw
```

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]
[.api-collapsible-fifth-title]
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per
unit of time.
```

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data".

"Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value.

"Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
```

```
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp
```

```
|string
```

a|The timestamp of the performance data.

```
|===
```

```
[#status]
```

```
[.api-collapsible-fifth-title]
```

status

Status information about the LUN.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

|Description

|container_state

|string

a|The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available.

|mapped

|boolean

a|Reports if the LUN is mapped to one or more initiator groups.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|read_only

|boolean

a|Reports if the LUN allows only read access.

|state

|string

a|The state of the LUN. Normal states for a LUN are `_online_` and `_offline_`. Other states indicate errors.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

The SVM in which the LUN is located.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the SVM.
```

```
|===
```

```
[#lun]
```

```
[.api-collapsible-fifth-title]
```

```
lun
```

A LUN is the logical representation of storage in a storage area network (SAN).

In ONTAP, a LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, and moved to a different volume. LUNs support the assignment of a quality of service (QoS) policy for performance management or a QoS policy can be assigned to the volume containing the LUN. See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the Fibre Channel Protocol or a TCP/IP network using iSCSI.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

a|

|auto_delete

|boolean

a|This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.

When set to `_true_`, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new LUN is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|class

|string

a|The class of LUN. Only `_regular_` LUNs can be created using the REST API.

|comment

|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time

|string

a|The time the LUN was created.

|enabled

|boolean

a|The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, you can consult the ``state`` property to determine if the LUN is administratively disabled (`_offline_`) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the ``enabled`` property to `_true_` or brought administratively offline by setting the ``enabled`` property to `_false_`. Upon creation, a LUN

is enabled by default. Valid in PATCH.

```
|location
|link:#location[location]
a|The location of the LUN within the ONTAP cluster. Valid in POST and
PATCH.
```

* Introduced in: 9.6

```
|lun_maps
|array[link:#lun_maps[lun_maps]]
a|The LUN maps with which the LUN is associated.
```

There is an added cost to retrieving property values for ``lun_maps``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|metric
|link:#metric[metric]
a|Performance numbers, such as IOPS latency and throughput.
```

```
|movement
|link:#movement[movement]
a|This sub-object applies to LUN movement between volumes. A LUN can be
moved to a new volume with a PATCH request that changes either the volume
portion of property `name`, `location.volume.uuid`, or
`location.volume.name`. If the volume is changed using more than one of
these properties, the supplied properties used must refer to the same
volume.
```

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The ``movement`` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the ``movement`` properties. The LUN movement operation can be further modified using a PATCH on the properties on the ``movement`` sub-object.

There is added cost to retrieving property values for `movement`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|name

|string

a|The fully qualified path name of the LUN composed of a `"/vol"` prefix, the volume name, the (optional) qtree name, and base name of the LUN. Valid in POST and PATCH.

A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.

A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.

|os_type

|string

a|The operating system type of the LUN.

Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone.

|qos_policy

|link:#qos_policy[qos_policy]

a|The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property ``qos_policy.uuid`` and ``qos_policy.name`` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property ``qos_policy.name`` to an empty string (`""`) in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

|serial_number

|string

a|The LUN serial number. The serial number is generated by ONTAP when the LUN is created.

```
* maxLength: 12
* minLength: 12
* readOnly: 1
* Introduced in: 9.6
```

```
|space
|link:#space[space]
a|The storage space related properties of the LUN.
```

```
|statistics
|link:#statistics[statistics]
a|These are raw performance numbers, such as IOPS latency and throughput.
These numbers are aggregated across all nodes in the cluster and increase
with the uptime of the cluster.
```

```
|status
|link:#status[status]
a|Status information about the LUN.
```

```
|svm
|link:#svm[svm]
a|The SVM in which the LUN is located.
```

```
|uuid
|string
a|The unique identifier of the LUN. The UUID is generated by ONTAP when
the LUN is created.
```

```
* example: 1cd8a442-86d1-11e0-ae1c-123478563412
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

|===

//end collapsible .Definitions block

====

[[ID421a6292c5b2945dbee8a566d96cb109]]

= Create a LUN

[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/storage/luns`#

Introduced In: 9.6

Creates a LUN.

== Required properties

* `svm.uuid` or `svm.name` - Existing SVM in which to create the LUN.

* `name`, `location.volume.name` or `location.volume.uuid` - Existing volume in which to create the LUN.

* `name` or `location.logical_unit` - Base name of the LUN.

* `os_type` - Operating system from which the LUN will be accessed.

Required when creating a non-clone LUN and disallowed when creating a clone of an existing LUN. A clone's `os_type` is taken from the source LUN.

* `space.size` - Size of the LUN. Required when creating a non-clone LUN and disallowed when creating a clone of an existing LUN. A clone's size is taken from the source LUN.

== Recommended optional properties

* `qos_policy.name` or `qos_policy.uuid` - Existing traditional or adaptive QoS policy to be applied to the LUN. All LUNs should be managed by a QoS policy at the volume or LUN level.

== Default property values

If not specified in POST, the follow default property values are assigned.

* `auto_delete` - `__false__`

== Related ONTAP commands

* `lun create`

* `volume file clone autodelete`

* `volume file clone create`

== Learn more

* xref:{relative_path}storage_luns_endpoint_overview.html[DOC
/storage/luns]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|auto_delete

|boolean

a|This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.

When set to `_true_`, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume

containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new LUN is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|class
|string
a|The class of LUN. Only _regular_ LUNs can be created using the REST API.
```

```
|clone
|link:#clone[clone]
a|This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, and `space.guarantee.requested`.
```

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: ``class``, ``auto_delete``, ``lun_maps``, ``serial_number``, ``status.state``, and ``uuid``.

Persistent reservations for the patched LUN are also preserved.

```
|comment
|string
a|A configurable comment available for use by the administrator. Valid in POST and PATCH.
```

```
|create_time
|string
a|The time the LUN was created.
```

```
|location
|link:#location[location]
```

a|The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

* Introduced in: 9.6

|lun_maps

|array[link:#lun_maps[lun_maps]]

a|The LUN maps with which the LUN is associated.

There is an added cost to retrieving property values for `lun_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|movement

|link:#movement[movement]

a|This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.

There is added cost to retrieving property values for `movement`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|name

|string

a|The fully qualified path name of the LUN composed of a "/vol" prefix, the volume name, the (optional) qtree name, and base name of the LUN. Valid in POST and PATCH.

A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.

A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.

|os_type

|string

a|The operating system type of the LUN.

Required in POST when creating a LUN that is not a clone of another.

Disallowed in POST when creating a LUN clone.

|qos_policy

|link:#qos_policy[qos_policy]

a|The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

|serial_number

|string

a|The LUN serial number. The serial number is generated by ONTAP when the LUN is created.

* maxLength: 12

* minLength: 12

* readOnly: 1

* Introduced in: 9.6

|space

|link:#space[space]

a|The storage space related properties of the LUN.

```
|status
|link:#status[status]
a|Status information about the LUN.
```

```
|svm
|link:#svm[svm]
a|The SVM in which the LUN is located.
```

```
|uuid
|string
a|The unique identifier of the LUN. The UUID is generated by ONTAP when
the LUN is created.
```

```
* example: 1cd8a442-86d1-11e0-ae1c-123478563412
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

.Example request

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
{
  "class": "string",
  "clone": {
    "source": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "create_time": "2018-06-04T19:00:00Z",
  "enabled": null,
  "location": {
    "logical_unit": "lun1",
    "qtree": {
      "id": "1",
      "name": "qt1"
    },
    "volume": {
```

```

    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "lun_maps": [
    {
      "igroup": {
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "logical_unit_number": 0
    }
  ],
  "movement": {
    "max_throughput": "string",
    "paths": {
      "destination": "/vol/vol1/lun1",
      "source": "/vol/vol2/lun2"
    }
  },
  "name": "/vol/volume1/mtree1/lun1",
  "os_type": "string",
  "qos_policy": {
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "serial_number": "string",
  "space": {
    "size": "1073741824",
    "used": 0
  },
  "status": {
    "container_state": "string",
    "state": "online"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response

```

Status: 201, Created

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|num_records
```

```
|integer
```

```
a|Number of records.
```

```
|records
```

```
|array[link:#lun[lun]]
```

```
a|
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "records": [
```

```
    {
```

```
      "class": "string",
```

```
      "clone": {
```

```
        "source": {
```

```
          "name": "/vol/volume1/lun1",
```

```
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
```

```
        }
```

```
      },
```

```
      "comment": "string",
```

```
      "create_time": "2018-06-04T19:00:00Z",
```

```
      "enabled": null,
```

```
      "location": {
```

```
        "logical_unit": "lun1",
```

```
        "qtree": {
```

```
          "id": "1",
```

```
          "name": "qt1"
```

```
        },
```

```
        "volume": {
```

```
          "name": "volume1",
```

```
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
```

```
        }
```

```
      },
```

```

    "lun_maps": [
      {
        "igroup": {
          "name": "igroup1",
          "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
        },
        "logical_unit_number": 0
      }
    ],
    "movement": {
      "max_throughput": "string",
      "paths": {
        "destination": "/vol/vol1/lun1",
        "source": "/vol/vol2/lun2"
      }
    },
    "name": "/vol/volume1/qtreen1/lun1",
    "os_type": "string",
    "qos_policy": {
      "name": "qos1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "serial_number": "string",
    "space": {
      "size": "1073741824",
      "used": 0
    },
    "status": {
      "container_state": "string",
      "state": "online"
    },
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

|===

| Error Code | Description

| 917927

| The specified volume was not found.

| 918236

| The specified `location.volume.uuid` and `location.volume.name` do not refer to the same volume.

| 2621462

| The specified SVM does not exist.

| 2621706

| The specified `svm.uuid` and `svm.name` do not refer to the same SVM.

| 2621707

| No SVM was specified. Either `svm.name` or `svm.uuid` must be supplied.

| 5242927

| The specified qtree was not found.

| 5242950

| The specified `location.qtree.id` and `location.qtree.name` do not refer to the same qtree.

| 5374121

| A LUN name can only contain characters A-Z, a-z, 0-9, "-", ".", "_", "{" and "}".

| 5374123

| A negative size was provided for the LUN.

| 5374124

| The specified size is too small for the LUN.

| 5374125

| The specified size is too large for the LUN.

| 5374129

| LUNs cannot be created on a load sharing mirror volume.

| 5374130

| An invalid size value was provided.

| 5374237
| LUNs cannot be created on an SVM root volume.

| 5374238
| LUNs cannot be created in Snapshot copies.

| 5374241
| A size value with invalid units was provided.

| 5374242
| A LUN or NVMe namespace already exists at the specified path.

| 5374352
| An invalid name was provided for the LUN.

| 5374707
| Creating a LUN in the specific volume is not allowed because the volume is reserved for an application.

| 5374858
| The volume specified by `name` is not the same as that specified by `location.volume`.

| 5374859
| No volume was specified for the LUN.

| 5374860
| The qtree specified by `name` is not the same as that specified by `location.qtree`.

| 5374861
| The LUN base name specified by `name` is not the same as that specified by `location.logical_unit`.

| 5374862
| No LUN path base name was provided for the LUN.

| 5374863
| An error occurred after successfully creating the LUN. Some properties were not set.

| 5374874
| The specified `clone.source.uuid` and `clone.source.name` do not refer to the same LUN.

| 5374875
| The specified `clone.source` was not found.

```

| 5374876
| The specified `clone.source` was not found.

| 5374883
| The property cannot be specified when creating a LUN clone. The `target`
property of the error object identifies the property.

| 5374884
| The property is required except when creating a LUN clone. The `target`
property of the error object identifies the property.

| 5374886
| An error occurred after successfully creating the LUN preventing the
retrieval of its properties.

| 5374899
| The `clone.source.uuid` property is not supported when specifying a
source LUN from a Snapshot copy.

| 13565952
| The LUN clone request failed.
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====

```

```

[#href]
[.api-collapsible-fifth-title]
href

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|href
|string
a|

```

```

|===

```



```
[#_links]
[.api-collapsible-fifth-title]
_links
[#source]
[.api-collapsible-fifth-title]
source
```

The source LUN for a LUN clone operation. This can be specified using property ``clone.source.uuid`` or ``clone.source.name``. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The fully qualified path name of the clone source LUN composed of a
"/vol" prefix, the volume name, the (optional) qtree name, and base name
of the LUN. Valid in POST and PATCH.
```

```
|uuid
|string
a|The unique identifier of the clone source LUN. Valid in POST and PATCH.
```

```
|===
```

```
[#clone]
[.api-collapsible-fifth-title]
clone
```

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN:

``auto_delete`, `qos_policy`, and `space.guarantee.requested`.`

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: ``class``, ``auto_delete``, ``lun_maps``, ``serial_number``, ``status.state``, and ``uuid``.

Persistent reservations for the patched LUN are also preserved.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|source
```

```
|link:#source[source]
```

a|The source LUN for a LUN clone operation. This can be specified using property ``clone.source.uuid`` or ``clone.source.name``. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

```
|===
```

```
[#qtree]
```

```
[.api-collapsible-fifth-title]
```

qtree

The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties ``name`` and ``location.qtree.name`` and/or ``location.qtree.uuid`` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
|Description
```

```
|id
|integer
```

```
a|The identifier for the qtree, unique within the qtree's volume.
```

```
|name
|string
```

```
a|The name of the qtree.
```

```
|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

The volume in which the LUN is located. Valid in POST and PATCH.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
```

```
a|The name of the volume.
```

```
|uuid
|string
```

```
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

* example: 028baa66-41bd-11e9-81d5-00a0986138f7

* Introduced in: 9.6

|===

[#location]

[.api-collapsible-fifth-title]

location

The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|logical_unit

|string

a|The base name component of the LUN. Valid in POST and PATCH.

If properties `name` and `location.logical_unit` are specified in the same request, they must refer to the base name.

A PATCH that modifies the base name of the LUN is considered a rename operation.

|qtree

|link:#qtree[qtree]

a|The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

|volume

|link:#volume[volume]

a|The volume in which the LUN is located. Valid in POST and PATCH.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

|===

```
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

The initiator group to which the LUN is mapped.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the initiator group.
```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

|===

```
[#lun_maps]
[.api-collapsible-fifth-title]
lun_maps
```

A LUN map with which the LUN is associated.

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description

|igroup
|link:#igroup[igroup]
a|The initiator group to which the LUN is mapped.

|logical_unit_number
|integer
a|The logical unit number assigned to the LUN for initiators in the
initiator group.
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
```

throughput

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

```
a|Performance metric for other I/O operations. Other I/O operations can be  
metadata operations, such as directory lookups and so on.
```

```
|read
```

```
|integer
```

```
a|Performance metric for read I/O operations.
```

```
|total
```

```
|integer
```

```
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
```

```
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

```
metric
```

Performance numbers, such as IOPS latency and throughput.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```



```

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#paths]

```

[.api-collapsible-fifth-title]

paths

The fully qualified LUN path names involved in the LUN movement.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|destination

|string

a|The fully qualified path of the LUN movement destination composed of a "/vol" prefix, the volume name, the (optional) qtree name, and base name of the LUN.

|source

|string

a|The fully qualified path of the LUN movement source composed of a "/vol" prefix, the volume name, the (optional) qtree name, and base name of the LUN.

|===

[#failure]

[.api-collapsible-fifth-title]

failure

Error information provided if the asynchronous LUN movement operation fails.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|The error code.

```
|message
|string
a|The error message.
```

```
|===
```

```
[#progress]
[.api-collapsible-fifth-title]
progress
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|elapsed
```

```
|integer
```

```
a|The amount of time, in seconds, that has elapsed since the start of the
LUN movement.
```

```
|failure
```

```
|link:#failure[failure]
```

```
a|Error information provided if the asynchronous LUN movement operation
fails.
```

```
|percent_complete
```

```
|integer
```

```
a|The percentage complete of the LUN movement.
```

```
|volume_snapshot_blocked
```

```
|boolean
```

```
a|This property reports if volume Snapshot copies are blocked by the LUN
movement. This property can be polled to identify when volume Snapshot
copies can be resumed after beginning a LUN movement.
```

```
|===
```

```
[#movement]
```

[.api-collapsible-fifth-title]

movement

This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property ``name``, ``location.volume.uuid``, or ``location.volume.name``. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The ``movement`` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the ``movement`` properties. The LUN movement operation can be further modified using a PATCH on the properties on the ``movement`` sub-object.

There is added cost to retrieving property values for ``movement``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|max_throughput

|string

a|The maximum data throughput that should be utilized in support of the LUN movement. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.

If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.

For more information, see `_Size properties_` in the `_docs_` section of the

ONTAP REST API documentation.

This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.

* Introduced in: 9.6

```
|paths
|link:#paths[paths]
a|The fully qualified LUN path names involved in the LUN movement.
```

```
|===
```

```
[#qos_policy]
[.api-collapsible-fifth-title]
qos_policy
```

The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property ``qos_policy.uuid`` and ``qos_policy.name`` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property ``qos_policy.name`` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set this property to an empty string ("" ) in a PATCH request. Valid in POST and PATCH.
```

```
|uuid
|string
a|The unique identifier of the QoS policy. Valid in POST and PATCH.
```

|===

```
[#guarantee]
[.api-collapsible-fifth-title]
guarantee
```

Properties that request and report the space guarantee for the LUN.

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|requested
```

```
|boolean
```

a|The requested space reservation policy for the LUN. If `_true_`, a space reservation is requested for the LUN; if `_false_`, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that the fractional reserve for the volume is 100%. Valid in POST and PATCH.

```
|reserved
```

```
|boolean
```

a|Reports if the LUN is space guaranteed.

If `_true_`, a space guarantee is requested and the containing volume and aggregate support the request. If `_false_`, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request.

|===

```
[#space]
[.api-collapsible-fifth-title]
space
```

The storage space related properties of the LUN.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|guarantee
```

```
|link:#guarantee[guarantee]
```

```
a|Properties that request and report the space guarantee for the LUN.
```

```
|size
```

```
|integer
```

```
a|The total provisioned size of the LUN. The LUN size can be increased but not be made smaller using the REST interface.
```

The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform and the available space in the containing volume and aggregate.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

```
* example: 1073741824
```

```
* Max value: 140737488355328
```

```
* Min value: 4096
```

```
* Introduced in: 9.6
```

```
|used
```

```
|integer
```

```
a|The amount of space consumed by the main data stream of the LUN.
```

This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

```
* readOnly: 1
```

```
* Introduced in: 9.6
```

```
|===
```

```
[#iops_raw]  
[.api-collapsible-fifth-title]  
iops_raw
```

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#latency_raw]  
[.api-collapsible-fifth-title]  
latency_raw
```


The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#throughput_raw]
```

```
[.api-collapsible-fifth-title]
```

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```

|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#statistics]
[.api-collapsible-fifth-title]
statistics

These are raw performance numbers, such as IOPS latency and throughput.
These numbers are aggregated across all nodes in the cluster and increase
with the uptime of the cluster.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|iops_raw
|link:#iops_raw[iops_raw]
a|The number of I/O operations observed at the storage object. This can be
used along with delta time to calculate the rate of I/O operations per

```

unit of time.

|latency_raw

|link:#latency_raw[latency_raw]

a|The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

|status

|string

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data".

"Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value.

"Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

|throughput_raw

|link:#throughput_raw[throughput_raw]

a|Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#status]

[.api-collapsible-fifth-title]

status

Status information about the LUN.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|container_state
```

```
|string
```

a|The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available.

```
|mapped
```

```
|boolean
```

a|Reports if the LUN is mapped to one or more initiator groups.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|read_only
```

```
|boolean
```

a|Reports if the LUN allows only read access.

```
|state
```

```
|string
```

a|The state of the LUN. Normal states for a LUN are `_online_` and `_offline_`. Other states indicate errors.

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

The SVM in which the LUN is located.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#lun]
[.api-collapsible-fifth-title]
lun
```

A LUN is the logical representation of storage in a storage area network (SAN).

In ONTAP, a LUN is located within a volume. Optionally, it can be located within a qtree in a volume.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, and moved to a different volume. LUNs support the assignment of a quality of service (QoS) policy for performance management or a QoS policy can be assigned to the volume containing the LUN. See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the Fibre Channel Protocol or a TCP/IP network using iSCSI.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|auto_delete
|boolean
```

a|This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.

When set to `_true_`, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new LUN is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|class
|string

a|The class of LUN. Only `_regular_` LUNs can be created using the REST API.

|clone
|link:#clone[clone]

a|This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: ``auto_delete``, ``qos_policy``, and ``space.guarantee.requested``.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: ``class``, ``auto_delete``, ``lun_maps``, ``serial_number``, ``status.state``, and ``uuid``.

Persistent reservations for the patched LUN are also preserved.

|comment
|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time
|string

a|The time the LUN was created.

|location

|link:#location[location]

a|The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

* Introduced in: 9.6

|lun_maps

|array[link:#lun_maps[lun_maps]]

a|The LUN maps with which the LUN is associated.

There is an added cost to retrieving property values for ``lun_maps``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

|movement

|link:#movement[movement]

a|This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property ``name``, ``location.volume.uuid``, or ``location.volume.name``. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The ``movement`` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the ``movement`` properties. The LUN movement operation can be further modified using a PATCH on the properties on the ``movement`` sub-object.

There is added cost to retrieving property values for ``movement``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requestin](#)

g_specific_fields[Requesting specific fields] to learn more.

|name

|string

a|The fully qualified path name of the LUN composed of a "/vol" prefix, the volume name, the (optional) qtree name, and base name of the LUN. Valid in POST and PATCH.

A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.

A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.

|os_type

|string

a|The operating system type of the LUN.

Required in POST when creating a LUN that is not a clone of another. Disallowed in POST when creating a LUN clone.

|qos_policy

|link:#qos_policy[qos_policy]

a|The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

|serial_number

|string

a|The LUN serial number. The serial number is generated by ONTAP when the LUN is created.

* maxLength: 12

* minLength: 12

* readOnly: 1

* Introduced in: 9.6


```
|space
|link:#space[space]
a|The storage space related properties of the LUN.
```

```
|status
|link:#status[status]
a|Status information about the LUN.
```

```
|svm
|link:#svm[svm]
a|The SVM in which the LUN is located.
```

```
|uuid
|string
a|The unique identifier of the LUN. The UUID is generated by ONTAP when
the LUN is created.
```

```
* example: 1cd8a442-86d1-11e0-ae1c-123478563412
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDc5b2a690695c0a3bcc12edbdae73b00c]]
= Delete a LUN

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
```

```
block]#`/storage/luns/{uuid}`#
```

Introduced In: 9.6

Deletes a LUN.

== Related ONTAP commands

* `lun delete`

== Learn more

* xref:{relative_path}storage_luns_endpoint_overview.html[DOC
/storage/luns]

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|uuid  
|string  
|path  
|True  
a|The unique identifier of the LUN.
```

```
|allow_delete_while_mapped  
|boolean  
|query  
|False
```

a|Allows deletion of a mapped LUN.

A mapped LUN might be in use. Deleting a mapped LUN also deletes the LUN map and makes the data no longer available. This might cause a disruption in the availability of data.

This parameter should be used with caution.

* Default value:

```
|===
```

== Response

Status: 200, Ok

== Error

Status: Default

ONTAP Error Response Codes

|===

Error Code	Description
------------	-------------

1254197	
---------	--

	The LUN is mapped and cannot be deleted without specifying the `allow_delete_while_mapped` query parameter.
--	-------------------------------------------------------------------------------------------------------------

5374705	
---------	--

	Deleting the LUN is not allowed because it is part of an application.
--	-----------------------------------------------------------------------

5374865	
---------	--

	The LUN's aggregate is offline. The aggregate must be online to modify or remove the LUN.
--	-------------------------------------------------------------------------------------------

5374866	
---------	--

	The LUN's volume is offline. The volume must be online to modify or remove the LUN.
--	-------------------------------------------------------------------------------------

5374875	
---------	--

	The specified LUN was not found.
--	----------------------------------

5374876	
---------	--

	The specified LUN was not found.
--	----------------------------------

|===

[cols=3*,options=header]

|===

Name

Type

Description

|error

```

|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID057d800bbc153cd6f216e7d4d4ef0c2d]]

```

= Retrieve LUN properties or data

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/storage/luns/{uuid}`#

Introduced In: 9.6

Retrieves a LUN.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

- * `auto_delete`
- * `lun_maps.+++`
- * `movement.+++`
- * `status.mapped`
- * `statistics.+++`
- * `metric.+++`

== Related ONTAP commands

- * `lun mapping show`
- * `lun move show`
- * `lun show`
- * `volume file clone show-autodelete`

== Learn more

- * [xref:{relative_path}storage_luns_endpoint_overview.html\[DOC /storage/luns\]](#)

== Parameters

[cols=5*,options=header]
|===

Name
Type
In
Required
Description

```
|uuid
|string
|path
|True
a|The unique identifier of the LUN to retrieve.
```

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|auto_delete
```

```
|boolean
```

a|This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.

When set to `_true_`, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new LUN is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See

xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields[Requesting specific fields] to learn more.

|class

|string

a|The class of LUN. Only `_regular_` LUNs can be created using the REST API.

|comment

|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time

|string

a|The time the LUN was created.

|enabled

|boolean

a|The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, you can consult the ``state`` property to determine if the LUN is administratively disabled (`_offline_`) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the ``enabled`` property to `_true_` or brought administratively offline by setting the ``enabled`` property to `_false_`. Upon creation, a LUN is enabled by default. Valid in PATCH.

|location

|link:#location[location]

a|The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

* Introduced in: 9.6

|lun_maps

|array[link:#lun_maps[lun_maps]]

a|The LUN maps with which the LUN is associated.

There is an added cost to retrieving property values for ``lun_maps``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See

xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields[Requesting specific fields] to learn more.

|metric

|link:#metric[metric]

a|Performance numbers, such as IOPS latency and throughput.

|movement

|link:#movement[movement]

a|This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property ``name``, ``location.volume.uuid``, or ``location.volume.name``. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The ``movement`` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the ``movement`` properties. The LUN movement operation can be further modified using a PATCH on the properties on the ``movement`` sub-object.

There is added cost to retrieving property values for ``movement``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields[Requesting specific fields] to learn more.

|name

|string

a|The fully qualified path name of the LUN composed of a `"/vol"` prefix, the volume name, the (optional) qtree name, and base name of the LUN. Valid in POST and PATCH.

A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.

A PATCH that modifies the volume portion of the LUN path begins an

asynchronous LUN movement operation.

|os_type
|string
a|The operating system type of the LUN.

Required in POST when creating a LUN that is not a clone of another.
Disallowed in POST when creating a LUN clone.

|qos_policy
|link:#qos_policy[qos_policy]
a|The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

|serial_number
|string
a|The LUN serial number. The serial number is generated by ONTAP when the LUN is created.

- * maxLength: 12
- * minLength: 12
- * readOnly: 1
- * Introduced in: 9.6

|space
|link:#space[space]
a|The storage space related properties of the LUN.

|statistics
|link:#statistics[statistics]
a|These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```

|status
|link:#status[status]
a|Status information about the LUN.

|svm
|link:#svm[svm]
a|The SVM in which the LUN is located.

|uuid
|string
a|The unique identifier of the LUN. The UUID is generated by ONTAP when
the LUN is created.

* example: 1cd8a442-86d1-11e0-ae1c-123478563412
* readOnly: 1
* Introduced in: 9.6

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "class": "string",
  "comment": "string",
  "create_time": "2018-06-04T19:00:00Z",
  "location": {
    "logical_unit": "lun1",
    "qtree": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": "1",
      "name": "qt1"
    }
  }
}

```

```

    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "lun_maps": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "igroup": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-aelc-123478563412"
      },
      "logical_unit_number": 0
    }
  ],
  "metric": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "duration": "PT15S",
    "iops": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency": {
      "read": "200",
      "total": "1000",
      "write": "100"
    }
  }
}

```

```

    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"movement": {
    "max_throughput": "string",
    "paths": {
        "destination": "/vol/vol1/lun1",
        "source": "/vol/vol2/lun2"
    },
    "progress": {
        "elapsed": 0,
        "failure": {
            "code": "4",
            "message": "Destination volume is offline."
        },
        "percent_complete": 0,
        "state": "string"
    }
},
"name": "/vol/volume1/qtree1/lun1",
"os_type": "string",
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"serial_number": "string",
"space": {
    "size": "1073741824",
    "used": 0
},
"statistics": {
    "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    }
}

```

```

    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"status": {
    "container_state": "string",
    "state": "online"
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 5374875
| The specified LUN was not found.
| 5374876
| The specified LUN was not found.

```

```

|===

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

```



```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]
a|

|===
```

```
[#source]
[.api-collapsible-fifth-title]
source
```

The source LUN for a LUN clone operation. This can be specified using property ``clone.source.uuid`` or ``clone.source.name``. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

```
[#clone]
[.api-collapsible-fifth-title]
clone
```

This sub-object is used in POST to create a new LUN as a clone of an

existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, and `space.guarantee.requested`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

```
[#qtree]
[.api-collapsible-fifth-title]
qtree
```

The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|id
|integer
a|The identifier for the qtree, unique within the qtree's volume.
```

```
|name
|string
a|The name of the qtree.
```

```
|===
```

```
[#volume]
[.api-collapsible-fifth-title]
volume
```

The volume in which the LUN is located. Valid in POST and PATCH.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the volume.
```

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|logical_unit
```

```
|string
```

a|The base name component of the LUN. Valid in POST and PATCH.

If properties ``name`` and ``location.logical_unit`` are specified in the same request, they must refer to the base name.

A PATCH that modifies the base name of the LUN is considered a rename operation.

```
|qtree
```

```
|link:#qtree[qtree]
```

a|The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties ``name`` and ``location.qtree.name`` and/or ``location.qtree.uuid`` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

```
|volume
```

```
|link:#volume[volume]
```

a|The volume in which the LUN is located. Valid in POST and PATCH.

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

```
|===
```

```
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

The initiator group to which the LUN is mapped.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the initiator group.
```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#lun_maps]
[.api-collapsible-fifth-title]
lun_maps
```

A LUN map with which the LUN is associated.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
```

a|

|igroup

|link:#igroup[igroup]

a|The initiator group to which the LUN is mapped.

|logical_unit_number

|integer

a|The logical unit number assigned to the LUN for initiators in the initiator group.

|===

[#iops]

[.api-collapsible-fifth-title]

iops

The rate of I/O operations observed at the storage object.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]  
[.api-collapsible-fifth-title]  
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|other  
|integer  
a|Performance metric for other I/O operations. Other I/O operations can be  
metadata operations, such as directory lookups and so on.
```

```
|read  
|integer  
a|Performance metric for read I/O operations.
```

```
|total  
|integer  
a|Performance metric aggregated over all types of I/O operations.
```

```
|write  
|integer  
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]  
[.api-collapsible-fifth-title]  
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#metric]
```

```
[.api-collapsible-fifth-title]
```

metric

Performance numbers, such as IOPS latency and throughput.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```



```
|link:#_links[_links]
```

```
a|
```

```
|duration
```

```
|string
```

```
a|The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
```

```
|iops
```

```
|link:#iops[iops]
```

```
a|The rate of I/O operations observed at the storage object.
```

```
|latency
```

```
|link:#latency[latency]
```

```
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
```

```
|string
```

```
a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
```

```
|throughput
```

```
|link:#throughput[throughput]
```

```
a|The rate of throughput bytes per second observed at the storage object.
```

```
|timestamp
```

```
|string
```

```
a|The timestamp of the performance data.
```

```
|===
```

```
[#paths]
[.api-collapsible-fifth-title]
paths
```

The fully qualified LUN path names involved in the LUN movement.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|destination
```

```
|string
```

```
a|The fully qualified path of the LUN movement destination composed of a
"/vol" prefix, the volume name, the (optional) qtree name, and base name
of the LUN.
```

```
|source
```

```
|string
```

```
a|The fully qualified path of the LUN movement source composed of a "/vol"
prefix, the volume name, the (optional) qtree name, and base name of the
LUN.
```

```
|===
```

```
[#failure]
```

```
[.api-collapsible-fifth-title]
```

```
failure
```

Error information provided if the asynchronous LUN movement operation fails.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

a|The error code.

|message

|string

a|The error message.

|===

[#progress]

[.api-collapsible-fifth-title]

progress

[cols=3*,options=header]

|===

|Name

|Type

|Description

|elapsed

|integer

a|The amount of time, in seconds, that has elapsed since the start of the LUN movement.

|failure

|link:#failure[failure]

a|Error information provided if the asynchronous LUN movement operation fails.

|percent_complete

|integer

a|The percentage complete of the LUN movement.

|state

|string

a|The state of the LUN movement.

Valid in PATCH when an LUN movement is active. Set to `_paused_` to pause a LUN movement. Set to `_replicating_` to resume a paused LUN movement.

|volume_snapshot_blocked

```
|boolean
a|This property reports if volume Snapshot copies are blocked by the LUN
movement. This property can be polled to identify when volume Snapshot
copies can be resumed after beginning a LUN movement.
```

```
|===
```

```
[#movement]
[.api-collapsible-fifth-title]
movement
```

This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property ``name``, ``location.volume.uuid``, or ``location.volume.name``. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The ``movement`` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the ``movement`` properties. The LUN movement operation can be further modified using a PATCH on the properties on the ``movement`` sub-object.

There is added cost to retrieving property values for ``movement``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|max_throughput
|string
```

a|The maximum data throughput that should be utilized in support of the LUN movement. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.

If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.

* Introduced in: 9.6

|paths

|link:#paths[paths]

a|The fully qualified LUN path names involved in the LUN movement.

|progress

|link:#progress[progress]

a|

|===

[#qos_policy]

[.api-collapsible-fifth-title]

qos_policy

The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property ``qos_policy.uuid`` and ``qos_policy.name`` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property ``qos_policy.name`` to an empty string (`""`) in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

[cols=3*,options=header]

|===

|Name

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the QoS policy. To remove the QoS policy from a LUN, leaving
it with no QoS policy, set this property to an empty string ("") in a
PATCH request. Valid in POST and PATCH.

|uuid
|string
a|The unique identifier of the QoS policy. Valid in POST and PATCH.

|===

[#guarantee]
[.api-collapsible-fifth-title]
guarantee

Properties that request and report the space guarantee for the LUN.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|requested
|boolean
a|The requested space reservation policy for the LUN. If _true_, a space
reservation is requested for the LUN; if _false_, the LUN is thin
provisioned. Guaranteeing a space reservation request for a LUN requires
that the volume in which the LUN resides is also space reserved and that
the fractional reserve for the volume is 100%. Valid in POST and PATCH.

|reserved
|boolean
a|Reports if the LUN is space guaranteed.

```

If `_true_`, a space guarantee is requested and the containing volume and aggregate support the request. If `_false_`, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request.

|===

[#space]
[.api-collapsible-fifth-title]
space

The storage space related properties of the LUN.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|guarantee

|link:#guarantee[guarantee]

a|Properties that request and report the space guarantee for the LUN.

|size

|integer

a|The total provisioned size of the LUN. The LUN size can be increased but not be made smaller using the REST interface.

The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform and the available space in the containing volume and aggregate.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* example: 1073741824

* Max value: 140737488355328

* Min value: 4096

* Introduced in: 9.6

|used

|integer

a|The amount of space consumed by the main data stream of the LUN.

This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* readOnly: 1

* Introduced in: 9.6

|===

[#iops_raw]

[.api-collapsible-fifth-title]

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.


```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
```

```
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]  
[.api-collapsible-fifth-title]  
throughput_raw
```

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|other  
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read  
|integer
```

a|Performance metric for read I/O operations.

```
|total  
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write  
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#statistics]  
[.api-collapsible-fifth-title]
```

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data".

"Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value.

"Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
```

```
|link:#throughput_raw[throughput_raw]
```

a|Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#status]
[.api-collapsible-fifth-title]
status
```

Status information about the LUN.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|container_state
|string
```

a|The state of the volume and aggregate that contain the LUN. LUNs are only available when their containers are available.

```
|mapped
|boolean
```

a|Reports if the LUN is mapped to one or more initiator groups.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|read_only
|boolean
```

a|Reports if the LUN allows only read access.

```
|state
|string
```

a|The state of the LUN. Normal states for a LUN are `_online_` and `_offline_`. Other states indicate errors.

|===

[#svm]

[.api-collapsible-fifth-title]

svm

The SVM in which the LUN is located.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the SVM.

|uuid

|string

a|The unique identifier of the SVM.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

```

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

```

====

[[ID8357d5137a41f2a4042e3e40f208f1a4]]

= Update an existing LUN

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-block]#`/storage/luns/{uuid}`#

Introduced In: 9.6

Updates the properties of a LUN. A PATCH request can also be used to overwrite the contents of a LUN as a clone of another, to begin movement of a LUN between volumes, and to pause and resume the movement of a LUN between volumes.

== Related ONTAP commands

- * `lun modify`
- * `lun move modify`
- * `lun move pause`
- * `lun move resume`
- * `lun move start`
- * `lun resize`
- * `volume file clone autodelete`

== Learn more

* xref:{relative_path}storage_luns_endpoint_overview.html[DOC /storage/luns]

== Parameters

[cols=5*,options=header]

|==

Name
Type
In
Required
Description

uuid
string
path

```
|True
a|The unique identifier of the LUN to update.
```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|auto_delete
```

```
|boolean
```

```
a|This property marks the LUN for auto deletion when the volume containing
the LUN runs out of space. This is most commonly set on LUN clones.
```

When set to `_true_`, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new LUN is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|class
```

```
|string
```

```
a|The class of LUN. Only _regular_ LUNs can be created using the REST API.
```

```
|clone
```

```
|link:#clone[clone]
```

```
a|This sub-object is used in POST to create a new LUN as a clone of an
existing LUN, or PATCH to overwrite an existing LUN as a clone of another.
Setting a property in this sub-object indicates that a LUN clone is
desired. Consider the following other properties when cloning a LUN:
`auto_delete`, `qos_policy`, and `space.guarantee.requested`.
```


When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: ``class``, ``auto_delete``, ``lun_maps``, ``serial_number``, ``status.state``, and ``uuid``.

Persistent reservations for the patched LUN are also preserved.

|comment

|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time

|string

a|The time the LUN was created.

|enabled

|boolean

a|The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, you can consult the ``state`` property to determine if the LUN is administratively disabled (`_offline_`) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the ``enabled`` property to `_true_` or brought administratively offline by setting the ``enabled`` property to `_false_`. Upon creation, a LUN is enabled by default. Valid in PATCH.

|location

|link:#location[location]

a|The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

* Introduced in: 9.6

|lun_maps

|array[link:#lun_maps[lun_maps]]

a|The LUN maps with which the LUN is associated.

There is an added cost to retrieving property values for ``lun_maps``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See

xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields[Requesting specific fields] to learn more.

|movement

|link:#movement[movement]

a|This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.

There is added cost to retrieving property values for `movement`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields[Requesting specific fields] to learn more.

|name

|string

a|The fully qualified path name of the LUN composed of a "/vol" prefix, the volume name, the (optional) qtree name, and base name of the LUN. Valid in POST and PATCH.

A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.

A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.

|qos_policy

|link:#qos_policy[qos_policy]

a|The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

|serial_number

|string

a|The LUN serial number. The serial number is generated by ONTAP when the LUN is created.

* maxLength: 12

* minLength: 12

* readOnly: 1

* Introduced in: 9.6

|space

|link:#space[space]

a|The storage space related properties of the LUN.

|status

|link:#status[status]

a|Status information about the LUN.

|svm

|link:#svm[svm]

a|The SVM in which the LUN is located.

|uuid

|string

a|The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.

* example: 1cd8a442-86d1-11e0-ae1c-123478563412

* readOnly: 1

* Introduced in: 9.6

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "class": "string",
  "clone": {
    "source": {
      "name": "/vol/volume1/lun1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "comment": "string",
  "create_time": "2018-06-04T19:00:00Z",
  "location": {
    "logical_unit": "lun1",
    "qtree": {
      "id": "1",
      "name": "qt1"
    },
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "lun_maps": [
    {
      "igroup": {
        "name": "igroup1",
        "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
      },
      "logical_unit_number": 0
    }
  ],
  "movement": {
    "max_throughput": "string",
    "paths": {
      "destination": "/vol/vol1/lun1",
      "source": "/vol/vol2/lun2"
    },
    "progress": {
      "elapsed": 0,
      "failure": {
```

```

        "code": "4",
        "message": "Destination volume is offline."
    },
    "percent_complete": 0,
    "state": "string"
}
},
"name": "/vol/volume1/mtree1/lun1",
"qos_policy": {
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"serial_number": "string",
"space": {
    "size": "1073741824",
    "used": 0
},
"status": {
    "container_state": "string",
    "state": "online"
},
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
====

== Response

```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```

|===
| Error Code | Description
| 917927
| The specified volume was not found.

```

| 918236
| The specified `location.volume.uuid` and `location.volume.name` do not refer to the same volume.

| 5242927
| The specified qtree was not found.

| 5242950
| The specified `location.qtree.id` and `location.qtree.name` do not refer to the same qtree.

| 5374124
| The specified LUN size is too small.

| 5374125
| The specified LUN size is too large.

| 5374130
| An invalid size value was provided.

| 5374241
| A size value with invalid units was provided.

| 5374480
| Modifying the LUN is not allowed because it is in a foreign LUN import relationship.

| 5374858
| The volume specified by `name` is not the same as that specified by `location.volume`.

| 5374860
| The qtree specified by `name` is not the same as that specified by `location.qtree`.

| 5374861
| The LUN base name specified by `name` is not the same as that specified by `location.logical_unit`.

| 5374864
| An error occurred after successfully overwriting data for the LUN as a clone. Some properties were not modified.

| 5374865
| The LUN's aggregate is offline. The aggregate must be online to modify or remove the LUN.

| 5374866
| The LUN's volume is offline. The volume must be online to modify or remove the LUN.

| 5374874
| The specified `clone.source.uuid` and `clone.source.name` do not refer to the same LUN.

| 5374875
| The specified LUN was not found. This can apply to `clone.source` or the target LUN. The `target` property of the error object identifies the property.

| 5374876
| The specified LUN was not found. This can apply to `clone.source` or the target LUN. The `target` property of the error object identifies the property.

| 5374885
| An error occurred after successfully modifying some of the properties of the LUN. Some properties were not modified.

| 5374889
| An invalid value was specified for `movement.progress.state`. Active LUN movement operations can be PATCHed to only `_paused_` or `_replicating_`.

| 5374892
| An attempt was made to reduce the size of a LUN.

| 5374904
| The destination volume is not online.

| 13565952
| The LUN clone request failed.
|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#source]
[.api-collapsible-fifth-title]
source
```

The source LUN for a LUN clone operation. This can be specified using property ``clone.source.uuid`` or ``clone.source.name``. If both properties are supplied, they must refer to the same LUN.

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The fully qualified path name of the clone source LUN composed of a
"/vol" prefix, the volume name, the (optional) qtree name, and base name
of the LUN. Valid in POST and PATCH.
```

```
|uuid
|string
a|The unique identifier of the clone source LUN. Valid in POST and PATCH.
```


|===

```
[#clone]
[.api-collapsible-fifth-title]
clone
```

This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: ``auto_delete``, ``qos_policy``, and ``space.guarantee.requested``.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: ``class``, ``auto_delete``, ``lun_maps``, ``serial_number``, ``status.state``, and ``uuid``.

Persistent reservations for the patched LUN are also preserved.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|source
|link:#source[source]
a|The source LUN for a LUN clone operation. This can be specified using
property `clone.source.uuid` or `clone.source.name`. If both properties
are supplied, they must refer to the same LUN.
```

Valid in POST to create a new LUN as a clone of the source.

Valid in PATCH to overwrite an existing LUN's data as a clone of another.

|===

```
[#qtree]
[.api-collapsible-fifth-title]
qtree
```

The qtree in which the LUN is optionally located. Valid in POST and PATCH.

If properties `name` and `location.qtree.name` and/or `location.qtree.uuid` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|id
```

```
|integer
```

a|The identifier for the qtree, unique within the qtree's volume.

```
|name
```

```
|string
```

a|The name of the qtree.

```
|===
```

```
[#volume]
```

```
[.api-collapsible-fifth-title]
```

volume

The volume in which the LUN is located. Valid in POST and PATCH.

If properties `name` and `location.volume.name` and/or `location.volume.uuid` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
|string
a|The name of the volume.
```

```
|uuid
|string
a|Unique identifier for the volume. This corresponds to the instance-uuid
that is exposed in the CLI and ONTAPI. It does not change due to a volume
move.
```

```
* example: 028baa66-41bd-11e9-81d5-00a0986138f7
* Introduced in: 9.6
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|logical_unit
|string
a|The base name component of the LUN. Valid in POST and PATCH.
```

If properties `name` and `location.logical_unit` are specified in the same request, they must refer to the base name.

A PATCH that modifies the base name of the LUN is considered a rename operation.

```
|qtree
|link:#qtree[qtree]
a|The qtree in which the LUN is optionally located. Valid in POST and
PATCH.
```

If properties ``name`` and ``location.qtree.name`` and/or ``location.qtree.uuid`` are specified in the same request, they must refer to the same qtree.

A PATCH that modifies the qtree of the LUN is considered a rename operation.

```
|volume
|link:#volume[volume]
a|The volume in which the LUN is located. Valid in POST and PATCH.
```

If properties ``name`` and ``location.volume.name`` and/or ``location.volume.uuid`` are specified in the same request, they must refer to the same volume.

A PATCH that modifies the volume of the LUN begins an asynchronous LUN movement operation.

```
|===
```

```
[#igroup]
[.api-collapsible-fifth-title]
igroup
```

The initiator group to which the LUN is mapped.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the initiator group.
```

```
|uuid
|string
a|The unique identifier of the initiator group.
```

```
|===
```

```
[#lun_maps]  
[.api-collapsible-fifth-title]  
lun_maps
```

A LUN map with which the LUN is associated.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|igroup
```

```
|link:#igroup[igroup]
```

a|The initiator group to which the LUN is mapped.

```
|logical_unit_number
```

```
|integer
```

a|The logical unit number assigned to the LUN for initiators in the initiator group.

```
|===
```

```
[#iops]  
[.api-collapsible-fifth-title]  
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency]
[.api-collapsible-fifth-title]
latency
```

The round trip latency in microseconds observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput]
[.api-collapsible-fifth-title]
throughput
```

The rate of throughput bytes per second observed at the storage object.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|other
```

```
|integer
```

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
```

```
|integer
```

a|Performance metric for read I/O operations.

```
|total
```

```
|integer
```

a|Performance metric aggregated over all types of I/O operations.

```
|write
```

```
|integer
```

a|Performance metric for write I/O operations.

```
|===
```

```
[#metric]
[.api-collapsible-fifth-title]
metric
```

Performance numbers, such as IOPS latency and throughput.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:
```

```
|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.
```

```
|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.
```

```
|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.
```

```
|throughput
|link:#throughput[throughput]
```


a|The rate of throughput bytes per second observed at the storage object.

|timestamp

|string

a|The timestamp of the performance data.

|===

[#paths]

[.api-collapsible-fifth-title]

paths

The fully qualified LUN path names involved in the LUN movement.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|destination

|string

a|The fully qualified path of the LUN movement destination composed of a "/vol" prefix, the volume name, the (optional) qtree name, and base name of the LUN.

|source

|string

a|The fully qualified path of the LUN movement source composed of a "/vol" prefix, the volume name, the (optional) qtree name, and base name of the LUN.

|===

[#failure]

[.api-collapsible-fifth-title]

failure

Error information provided if the asynchronous LUN movement operation fails.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|The error code.
```

```
|message
```

```
|string
```

```
a|The error message.
```

```
|===
```

```
[#progress]
```

```
[.api-collapsible-fifth-title]
```

```
progress
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|elapsed
```

```
|integer
```

```
a|The amount of time, in seconds, that has elapsed since the start of the LUN movement.
```

```
|failure
```

```
|link:#failure[failure]
```

```
a|Error information provided if the asynchronous LUN movement operation fails.
```

```
|percent_complete
```

```
|integer
```

```
a|The percentage complete of the LUN movement.
```

```
|state
|string
a|The state of the LUN movement.
```

Valid in PATCH when an LUN movement is active. Set to `_paused_` to pause a LUN movement. Set to `_replicating_` to resume a paused LUN movement.

```
|volume_snapshot_blocked
|boolean
a|This property reports if volume Snapshot copies are blocked by the LUN
movement. This property can be polled to identify when volume Snapshot
copies can be resumed after beginning a LUN movement.
```

```
|===
```

```
[#movement]
[.api-collapsible-fifth-title]
movement
```

This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property ``name``, ``location.volume.uuid``, or ``location.volume.name``. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The ``movement`` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the ``movement`` properties. The LUN movement operation can be further modified using a PATCH on the properties on the ``movement`` sub-object.

There is added cost to retrieving property values for ``movement``. They are not populated for either a collection GET or an instance GET unless explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|max_throughput
```

```
|string
```

a|The maximum data throughput that should be utilized in support of the LUN movement. This property can be used to throttle a transfer and limit its impact on the performance of the source and destination nodes. The specified value will be rounded up to the nearest megabyte.

If this property is not specified in a POST that begins a LUN movement, throttling is not applied to the data transfer.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

This property is valid only in a POST that begins a LUN movement or a PATCH when a LUN movement is already in process.

* Introduced in: 9.6

```
|paths
```

```
|link:#paths[paths]
```

a|The fully qualified LUN path names involved in the LUN movement.

```
|progress
```

```
|link:#progress[progress]
```

```
a|
```

```
|===
```

```
[#qos_policy]
```

```
[.api-collapsible-fifth-title]
```

```
qos_policy
```

The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property ``qos_policy.uuid`` and ``qos_policy.name`` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set

property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|name
```

```
|string
```

a|The name of the QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set this property to an empty string ("") in a PATCH request. Valid in POST and PATCH.

```
|uuid
```

```
|string
```

a|The unique identifier of the QoS policy. Valid in POST and PATCH.

```
|===
```

```
[#guarantee]
```

```
[.api-collapsible-fifth-title]
```

guarantee

Properties that request and report the space guarantee for the LUN.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|requested
```

```
|boolean
```

a|The requested space reservation policy for the LUN. If `_true_`, a space reservation is requested for the LUN; if `_false_`, the LUN is thin provisioned. Guaranteeing a space reservation request for a LUN requires that the volume in which the LUN resides is also space reserved and that

the fractional reserve for the volume is 100%. Valid in POST and PATCH.

|reserved
|boolean
a|Reports if the LUN is space guaranteed.

If `_true_`, a space guarantee is requested and the containing volume and aggregate support the request. If `_false_`, a space guarantee is not requested or a space guarantee is requested and either the containing volume or aggregate do not support the request.

|===

[#space]
[.api-collapsible-fifth-title]
space

The storage space related properties of the LUN.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|guarantee
|link:#guarantee[guarantee]
a|Properties that request and report the space guarantee for the LUN.

|size
|integer
a|The total provisioned size of the LUN. The LUN size can be increased but not be made smaller using the REST interface.

The maximum and minimum sizes listed here are the absolute maximum and absolute minimum sizes in bytes. The actual minimum and maximum sizes vary depending on the ONTAP version, ONTAP platform and the available space in the containing volume and aggregate.
For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* example: 1073741824

* Max value: 140737488355328

* Min value: 4096

* Introduced in: 9.6

|used

|integer

a|The amount of space consumed by the main data stream of the LUN.

This value is the total space consumed in the volume by the LUN, including filesystem overhead, but excluding prefix and suffix streams. Due to internal filesystem overhead and the many ways SAN filesystems and applications utilize blocks within a LUN, this value does not necessarily reflect actual consumption/availability from the perspective of the filesystem or application. Without specific knowledge of how the LUN blocks are utilized outside of ONTAP, this property should not be used as an indicator for an out-of-space condition.

For more information, see `_Size properties_` in the `_docs_` section of the ONTAP REST API documentation.

* readOnly: 1

* Introduced in: 9.6

|===

[#iops_raw]

[.api-collapsible-fifth-title]

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#latency_raw]
[.api-collapsible-fifth-title]
latency_raw
```

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
```



```
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#throughput_raw]
[.api-collapsible-fifth-title]
throughput_raw
```

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

```
|===
```

```
[#statistics]  
[.api-collapsible-fifth-title]  
statistics
```

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|iops_raw
```

```
|link:#iops_raw[iops_raw]
```

a|The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

```
|latency_raw
```

```
|link:#latency_raw[latency_raw]
```

a|The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

```
|status
```

```
|string
```

a|Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

```
|throughput_raw
|link:#throughput_raw[throughput_raw]
a|Throughput bytes observed at the storage object. This can be used along
with delta time to calculate the rate of throughput bytes per unit of
time.
```

```
|timestamp
|string
a|The timestamp of the performance data.
```

```
|===
```

```
[#status]
[.api-collapsible-fifth-title]
status
```

Status information about the LUN.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|container_state
|string
a|The state of the volume and aggregate that contain the LUN. LUNs are
only available when their containers are available.
```

```
|mapped
|boolean
a|Reports if the LUN is mapped to one or more initiator groups.
```

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|read_only
```

```
|boolean
a|Reports if the LUN allows only read access.
```

```
|state
|string
a|The state of the LUN. Normal states for a LUN are _online_ and
_offline_. Other states indicate errors.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

The SVM in which the LUN is located.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#lun]
[.api-collapsible-fifth-title]
lun
```

A LUN is the logical representation of storage in a storage area network (SAN).

In ONTAP, a LUN is located within a volume. Optionally, it can be located

within a qtree in a volume.

A LUN can be created to a specified size using thin or thick provisioning. A LUN can then be renamed, resized, cloned, and moved to a different volume. LUNs support the assignment of a quality of service (QoS) policy for performance management or a QoS policy can be assigned to the volume containing the LUN. See the LUN object model to learn more about each of the properties supported by the LUN REST API.

A LUN must be mapped to an initiator group to grant access to the initiator group's initiators (client hosts). Initiators can then access the LUN and perform I/O over a Fibre Channel (FC) fabric using the Fibre Channel Protocol or a TCP/IP network using iSCSI.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|auto_delete
```

```
|boolean
```

a|This property marks the LUN for auto deletion when the volume containing the LUN runs out of space. This is most commonly set on LUN clones.

When set to `_true_`, the LUN becomes eligible for automatic deletion when the volume runs out of space. Auto deletion only occurs when the volume containing the LUN is also configured for auto deletion and free space in the volume decreases below a particular threshold.

This property is optional in POST and PATCH. The default value for a new LUN is `_false_`.

There is an added cost to retrieving this property's value. It is not populated for either a collection GET or an instance GET unless it is explicitly requested using the ``fields`` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|class
```

```
|string
```

a|The class of LUN. Only `_regular_` LUNs can be created using the REST API.

```
|clone
```

|link:#clone[clone]

a|This sub-object is used in POST to create a new LUN as a clone of an existing LUN, or PATCH to overwrite an existing LUN as a clone of another. Setting a property in this sub-object indicates that a LUN clone is desired. Consider the following other properties when cloning a LUN: `auto_delete`, `qos_policy`, and `space.guarantee.requested`.

When used in a PATCH, the patched LUN's data is over-written as a clone of the source and the following properties are preserved from the patched LUN unless otherwise specified as part of the PATCH: `class`, `auto_delete`, `lun_maps`, `serial_number`, `status.state`, and `uuid`.

Persistent reservations for the patched LUN are also preserved.

|comment

|string

a|A configurable comment available for use by the administrator. Valid in POST and PATCH.

|create_time

|string

a|The time the LUN was created.

|enabled

|boolean

a|The enabled state of the LUN. LUNs can be disabled to prevent access to the LUN. Certain error conditions also cause the LUN to become disabled. If the LUN is disabled, you can consult the `state` property to determine if the LUN is administratively disabled (`_offline_`) or has become disabled as a result of an error. A LUN in an error condition can be brought online by setting the `enabled` property to `_true_` or brought administratively offline by setting the `enabled` property to `_false_`. Upon creation, a LUN is enabled by default. Valid in PATCH.

|location

|link:#location[location]

a|The location of the LUN within the ONTAP cluster. Valid in POST and PATCH.

* Introduced in: 9.6

|lun_maps

```
|array[link:#lun_maps[lun_maps]]
```

a|The LUN maps with which the LUN is associated.

There is an added cost to retrieving property values for `lun_maps`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|movement
```

```
|link:#movement[movement]
```

a|This sub-object applies to LUN movement between volumes. A LUN can be moved to a new volume with a PATCH request that changes either the volume portion of property `name`, `location.volume.uuid`, or `location.volume.name`. If the volume is changed using more than one of these properties, the supplied properties used must refer to the same volume.

Moving a LUN between volumes is an asynchronous activity begun by a PATCH request. The data for the LUN is then asynchronously copied from the source volume to the destination volume. The time required to complete the move depends on the size of the LUN and the load on the cluster. The `movement` sub-object is populated while a LUN movement is in progress and for two (2) minutes following completion of a movement.

While the LUN is being moved, the status of the LUN movement operation can be obtained using a GET for the LUN that requests the `movement` properties. The LUN movement operation can be further modified using a PATCH on the properties on the `movement` sub-object.

There is added cost to retrieving property values for `movement`. They are not populated for either a collection GET or an instance GET unless explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

```
|name
```

```
|string
```

a|The fully qualified path name of the LUN composed of a "/vol" prefix, the volume name, the (optional) qtree name, and base name of the LUN. Valid in POST and PATCH.

A PATCH that modifies the qtree and/or base name portion of the LUN path is considered a rename operation.

A PATCH that modifies the volume portion of the LUN path begins an asynchronous LUN movement operation.

|qos_policy

|link:#qos_policy[qos_policy]

a|The QoS policy for the LUN. Both traditional and adaptive QoS policies are supported. If both property `qos_policy.uuid` and `qos_policy.name` are specified in the same request, they must refer to the same QoS policy. To remove the QoS policy from a LUN, leaving it with no QoS policy, set property `qos_policy.name` to an empty string ("") in a PATCH request. Valid in POST and PATCH.

Note that a QoS policy can be set on a LUN, or a LUN's volume, but not both.

|serial_number

|string

a|The LUN serial number. The serial number is generated by ONTAP when the LUN is created.

* maxLength: 12

* minLength: 12

* readOnly: 1

* Introduced in: 9.6

|space

|link:#space[space]

a|The storage space related properties of the LUN.

|status

|link:#status[status]

a|Status information about the LUN.

|svm

|link:#svm[svm]

a|The SVM in which the LUN is located.

|uuid

|string

a|The unique identifier of the LUN. The UUID is generated by ONTAP when the LUN is created.


```
* example: 1cd8a442-86d1-11e0-ae1c-123478563412
* readOnly: 1
* Introduced in: 9.6
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID84ba5611af6b6f45b340340138fd00db]]
= Retrieve historical performance metrics for a LUN

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/storage/luns/{uuid}/metrics`#

*Introduced In:* 9.7

Retrieves historical performance metrics for a LUN.

== Parameters

[cols=5*,options=header]
|===

|Name
|Type
|In
|Required
|Description

|throughput.total
|integer

```

```
|query
|False
a|Filter by throughput.total

|throughput.read
|integer
|query
|False
a|Filter by throughput.read

|throughput.other
|integer
|query
|False
a|Filter by throughput.other

|throughput.write
|integer
|query
|False
a|Filter by throughput.write

|duration
|string
|query
|False
a|Filter by duration

|timestamp
|string
|query
|False
a|Filter by timestamp

|iops.total
|integer
|query
|False
a|Filter by iops.total
```

```
|iops.read
|integer
|query
|False
a|Filter by iops.read
```

```
|iops.other
|integer
|query
|False
a|Filter by iops.other
```

```
|iops.write
|integer
|query
|False
a|Filter by iops.write
```

```
|status
|string
|query
|False
a|Filter by status
```

```
|latency.total
|integer
|query
|False
a|Filter by latency.total
```

```
|latency.read
|integer
|query
|False
a|Filter by latency.read
```

```
|latency.other
|integer
|query
|False
a|Filter by latency.other
```

```
|latency.write
|integer
|query
|False
a|Filter by latency.write
```

```
|uuid
|string
|path
|True
a|Unique identifier of the LUN.
```

```
|interval
|string
|query
|False
a|The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y.
The period for each time range is as follows:
```

- * 1h: Metrics over the most recent hour sampled over 15 seconds.
- * 1d: Metrics over the most recent day sampled over 5 minutes.
- * 1w: Metrics over the most recent week sampled over 30 minutes.
- * 1m: Metrics over the most recent month sampled over 2 hours.
- * 1y: Metrics over the most recent year sampled over a day.
- * Default value: 1
- * enum: ["1h", "1d", "1w", "1m", "1y"]

```
|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.
```

- * Default value: 1
- * Max value: 120
- * Min value: 0

```
|fields
```

```
|array[string]
|query
|False
a|Specify the fields to return.

|max_records
|integer
|query
|False
a|Limit the number of records returned.

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|return_records
|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|===

== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|num_records
|integer
a|Number of records
```

```
|records
|array[link:#records[records]]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
```

```

        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
}
]
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

```



```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|self
|link:#href[href]
a|
```

```
|===
```

```
[#iops]
[.api-collapsible-fifth-title]
iops
```

The rate of I/O operations observed at the storage object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.
```

```
|read
|integer
a|Performance metric for read I/O operations.
```

```
|total
|integer
a|Performance metric aggregated over all types of I/O operations.
```

```
|write
|integer
a|Performance metric for write I/O operations.
```

|===

[#latency]
[.api-collapsible-fifth-title]
latency

The round trip latency in microseconds observed at the storage object.

[cols=3*,options=header]

|===

|Name
|Type
|Description

|other

|integer

a|Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

|read

|integer

a|Performance metric for read I/O operations.

|total

|integer

a|Performance metric aggregated over all types of I/O operations.

|write

|integer

a|Performance metric for write I/O operations.

|===

[#throughput]
[.api-collapsible-fifth-title]
throughput

The rate of throughput bytes per second observed at the storage object.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|other
|integer
a|Performance metric for other I/O operations. Other I/O operations can be
metadata operations, such as directory lookups and so on.

|read
|integer
a|Performance metric for read I/O operations.

|total
|integer
a|Performance metric aggregated over all types of I/O operations.

|write
|integer
a|Performance metric for write I/O operations.

|===

[#records]
[.api-collapsible-fifth-title]
records

Performance numbers, such as IOPS latency and throughput.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|duration
|string
a|The duration over which this sample is calculated. The time durations
are represented in the ISO-8601 standard format. Samples can be calculated
over the following durations:

|iops
|link:#iops[iops]
a|The rate of I/O operations observed at the storage object.

|latency
|link:#latency[latency]
a|The round trip latency in microseconds observed at the storage object.

|status
|string
a|Errors associated with the sample. For example, if the aggregation of
data over multiple nodes fails, then any partial errors might return "ok"
on success or "error" on an internal uncategorized failure. Whenever a
sample collection is missed but done at a later time, it is back filled to
the previous 15 second timestamp and tagged with "backfilled_data".
"Inconsistent_delta_time" is encountered when the time between two
collections is not the same for all nodes. Therefore, the aggregated value
might be over or under inflated. "Negative_delta" is returned when an
expected monotonically increasing value has decreased in value.
"Inconsistent_old_data" is returned when one or more nodes do not have the
latest data.

|throughput
|link:#throughput[throughput]
a|The rate of throughput bytes per second observed at the storage object.

|timestamp
|string
a|The timestamp of the performance data.

|===

[#error_arguments]

```

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

|message

|string

a|Message argument

|===

[#error]

[.api-collapsible-fifth-title]

error

[cols=3*,options=header]

|===

|Name

|Type

|Description

|arguments

|array[link:#error_arguments[error_arguments]]

a|Message arguments

|code

|string

a|Error code

|message

|string

a|Error message

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
:leveloffset: -1
```

```
:leveloffset: -1
```

```
= SVM
```

```
:leveloffset: +1
```

```
[[ID96bb38a8d9c792b4d88f22c7fd0422a5]]
= SVM overview
```

```
== Overview
```

Storage Virtual Machine (SVM) APIs enable you to manage SVMs and their attributes, including the configuration of the CIFS and NFS protocols, export policies, name mappings between CIFS and NFS users, and network services.

SVMs contain data volumes and one or more network interfaces through which they serve data to the clients. SVMs securely isolate the shared virtualized data storage and network, and each SVM appears as a single dedicated server to the clients. Each SVM has a separate administrator authentication domain, and each SVM can be managed independently by its SVM administrator.

In a cluster, SVMs facilitate data access. A cluster must have at least one SVM to serve data. Multiple SVMs can coexist in a single cluster without being bound to any particular node in the cluster. However, they are bound to the physical cluster on which they exist.

SVMs with volumes can contain files and LUNs. They provide file-level data access by using NFS and CIFS protocols for the NAS clients, and block-level data access by using iSCSI and Fibre Channel (FC) (FCoE included) for SAN hosts. The volumes within each NAS SVM are related to each other through junctions and they are mounted on junction paths. These junctions present the file system in each volume. The root volume of the SVM resides at the top level of the namespace hierarchy; additional volumes are mounted to the SVM root volume to extend the namespace. As volumes are created for the SVM, the root volume of the SVM contains junction paths.

The ONTAP REST APIs only expose data SVMs as an SVM. The information and configuration associated with the cluster and nodes are exposed from REST, but the cluster and nodes are not treated as if they are a type of SVM from REST APIs. Some APIs that expose both cluster-owned resources and SVM-owned resources from the same endpoint only return and support the "svm" sub-object for the resources that are within a data SVM. In those endpoints, the resources that are not in a data SVM do not return the "svm" sub-object. Generally, such endpoints have a "scope" attribute that returns either "svm" or "cluster" to identify the resource as either a cluster-level resource or one that is completely contained in a data SVM.

= Manage SVM peer permissions

:leveloffset: +1

[[IDe2cbb316ae8a0691b49d39de8b48e7ba]]

= SVM peer-permissions endpoint overview

== Managing SVM peer permissions

A cluster administrator can provide permissions for use during intercluster SVM peer relationship creation. Once this permission exists for a local SVM and peer cluster combination on a local cluster, no explicit SVM peer accept (or REST PATCH) API is required for any incoming SVM peer relationship creation requests from a remote cluster for that local SVM. Peer relationship directly changes the state to peered on both clusters. Use an SVM name as "*" to create permissions that apply to all local SVMs.

=== SVM peer permission APIs

The following APIs are used to manage SVM peer permissions:

- * GET /api/svm/peer-permissions
- * POST /api/svm/peer-permissions
- * GET /api/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}
- * PATCH /api/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}
- * DELETE /api/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}

[[ID762ffb4d80cbfdd44c991400d1ee0bb0]]

= Retrieve SVM peer permissions

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/svm/peer-permissions`#

Introduced In: 9.6

Retrieves the list of SVM peer permissions.

== Related ONTAP commands

* `vserver peer permission show`

== Examples

The following examples show how to retrieve a collection of SVM peer permissions based on a query.

. Retrieves a list of SVM peer permissions of a specific local SVM

GET "/api/svm/peer-permissions/?svm.name=VS1"

. Retrieves a list of SVM peer permissions of a specific cluster peer

GET "/api/svm/peer-permissions/?cluster_peer.name=cluster2"

== Learn more

```
* xref:{relative_path}svm_peer-permissions_endpoint_overview.html[DOC  
/svm/peer-permissions]
```

```
== Parameters
```

```
[cols=5*,options=header]
```

```
|==
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|cluster_peer.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by cluster_peer.name
```

```
|cluster_peer.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by cluster_peer.uuid
```

```
|applications
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by applications
```

```
|svm.uuid
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by svm.uuid
```

```
|svm.name
```

```
|string
```

```
|query
```

```
|False
```

a|Filter by svm.name

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|max_records

|integer

|query

|False

a|Limit the number of records returned.

|return_records

|boolean

|query

|False

a|The default is true for GET calls. When set to false, only the number of records is returned.

* Default value: 1

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|order_by

|array[string]

|query

|False

a|Order results by specified fields and optional [asc|desc] direction. Default direction is 'asc' for ascending.

```
|===
```

```
== Response
```

Status: 200, Ok

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|num_records
```

```
|integer
```

```
a|Number of records
```

```
|records
```

```
|array[link:#svm_peer_permission[svm_peer_permission]]
```

```
a|
```

```
|===
```

```
.Example response
```

```
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
```

```
  "_links": {
```

```
    "next": {
```

```
      "href": "/api/resourcelink"
```

```
    },
```

```
    "self": {
```

```
      "href": "/api/resourcelink"
```

```
    }
```

```
  },
```

```
  "records": [
```

```
    {
```

```
      "_links": {
```

```
        "self": {
```

```

        "href": "/api/resourcelink"
    }
},
"applications": [
    "snapmirror",
    "flexcache"
],
"cluster_peer": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```
|===
```

```
| Error codes | Description
```

```
| 26345574
```

```
| Failed to find the SVM or volume name with UUID.
```

```
|===
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

```

```
[#cluster_peer]
[.api-collapsible-fifth-title]
cluster_peer
```

Peer cluster details

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "*".

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
```



```

|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#svm_peer_permission]
[.api-collapsible-fifth-title]
svm_peer_permission

Manage SVM peer permissions.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|applications
|array[string]
a|A list of applications for an SVM peer relation.

|cluster_peer
|link:#cluster_peer[cluster_peer]
a|Peer cluster details

|svm
|link:#svm[svm]
a|Local SVM permitted for peer relation. To create peer permissions for
all SVMs, specify the SVM name as "*".

|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
=====
```

```
[[ID61a1e00346fd6243fadbdca6efb11ba4]]
= Create an SVM peer permission
```

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/svm/peer-permissions`#
```

```
*Introduced In:* 9.6
```

Creates an SVM peer permission.

== Required properties

- * `svm.name` or `svm.uuid` - SVM name
- * `cluster_peer.uuid` or `cluster_peer.name` - Peer cluster name or peer cluster UUID
- * `applications` - Peering applications

== Related ONTAP commands

- * `vserver peer permission create`

== Examples

The following examples show how to create SVM peer permissions.

. Creates an SVM peer permission entry with the local SVM and cluster peer names

```
-----
POST "/api/svm/peer-permissions" '{"cluster_peer":{"name":"cluster2"},
"svm":{"name":"VS1"}, "applications":["snapmirror"]}'
-----
```

. Creates an SVM peer permission entry with the local SVM and cluster peer UUID

```
POST "/api/svm/peer-permissions" '{"cluster_peer":{"uuid":"d3268a74-ee76-11e8-a9bb-005056ac6dc9"}, "svm":{"uuid":"8f467b93-f2f1-11e8-9027-005056ac81fc"}, "applications":["snapmirror"]}'
```

. Creates an SVM peer permission entry with all SVMs and the cluster peer name

```
POST "/api/svm/peer-permissions" '{"cluster_peer":{"name":"cluster2"}, "svm":{"name":"*"}, "applications":["snapmirror"]}'
```

== Learn more

* xref:{relative_path}svm_peer-permissions_endpoint_overview.html[DOC /svm/peer-permissions]

== Parameters

[cols=5*,options=header]
|==

Name
Type
In
Required
Description

return_records
boolean
query
False

a|The default is false. If set to true, the records are returned.

* Default value:

|==

== Request Body

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|applications
|array[string]
a|A list of applications for an SVM peer relation.


|cluster_peer
|link:#cluster_peer[cluster_peer]
a|Peer cluster details


|svm
|link:#svm[svm]
a|Local SVM permitted for peer relation. To create peer permissions for
all SVMs, specify the SVM name as "*".


|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebbf2"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

```

== Response

Status: 201, Created

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|applications
|array[string]
a|A list of applications for an SVM peer relation.

|cluster_peer
|link:#cluster_peer[cluster_peer]
a|Peer cluster details

|svm
|link:#svm[svm]
a|Local SVM permitted for peer relation. To create peer permissions for
all SVMs, specify the SVM name as "*".

|===

.Example response
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
=====

== Error

```

ONTAP Error Response Codes

|===

| Error codes | Description

| 26345572

| \{field\} is a required field.

| 26345573

| Failed to find the SVM or volume UUID with name.

| 26345574

| Failed to find the SVM or volume name with UUID.

| 26345575

| The specified peer cluster name and peer cluster UUID do not match.

|===

== Definitions

[.api-def-first-level]

.See Definitions

[%collapsible%closed]

//Start collapsible Definitions block

=====

[#href]

[.api-collapsible-fifth-title]

href

[cols=3*,options=header]

|===

|Name

|Type

|Description

|href

|string

a|

|===

[#_links]


```
[.api-collapsible-fifth-title]
_links
[#cluster_peer]
[.api-collapsible-fifth-title]
cluster_peer
```

Peer cluster details

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "*".

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
```

a|The unique identifier of the SVM.

|===

```
[#svm_peer_permission]
[.api-collapsible-fifth-title]
svm_peer_permission
```

Manage SVM peer permissions.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|applications
|array[string]
a|A list of applications for an SVM peer relation.
```

```
|cluster_peer
|link:#cluster_peer[cluster_peer]
a|Peer cluster details
```

```
|svm
|link:#svm[svm]
a|Local SVM permitted for peer relation. To create peer permissions for
all SVMs, specify the SVM name as "*".
```

|===

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====

[[ID779418c79d737be450a5832194e64b0d]]
= Delete SVM peer permissions

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}`#

*Introduced In:* 9.6

Deletes the SVM peer permissions.

== Related ONTAP commands

* `verver peer permission delete`

== Example

Deletes an SVM peer permission.

----
DELETE "/api/svm/peer-permissions/d3268a74-ee76-11e8-a9bb-
005056ac6dc9/8f467b93-f2f1-11e8-9027-005056ac81fc"
----

== Learn more

* xref:{relative_path}svm_peer-permissions_endpoint_overview.html[DOC
/svm/peer-permissions]

== Parameters

[cols=5*,options=header]
|===
|Name
|Type
|In
|Required
|Description

|cluster_peer.uuid
|string
```

```
|path
|True
a|Peer cluster UUID

* Introduced in: 9.7

|svm.uuid
|string
|path
|True
a|SVM UUID

|===

== Response
```

Status: 200, Ok

```
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error codes | Description

| 26345574
| Failed to find the SVM or volume name with UUID.
|===
```

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|error
|link:#error[error]
a|
```

```
|===
```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDdfbf8c7c7792c5912484635afcc9c654]]
= Retrieve an SVM peer permission instance

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}`#

```

Introduced In: 9.6

Retrieves the SVM peer permission instance.

== Related ONTAP commands

* `vserver peer permission show`

== Example

The following example shows how to retrieve the parameters for an SVM peer permission.

```
GET "/api/svm/peer-permissions/d3268a74-ee76-11e8-a9bb-005056ac6dc9/8f467b93-f2f1-11e8-9027-005056ac81fc"
```

== Learn more

* xref:{relative_path}svm_peer-permissions_endpoint_overview.html[DOC /svm/peer-permissions]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|cluster_peer.uuid

|string

|path

|True

a|Peer cluster UUID

* Introduced in: 9.7

|svm.uuid

|string


```

|path
|True
a|SVM UUID

|fields
|array[string]
|query
|False
a|Specify the fields to return.

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|applications
|array[string]
a|A list of applications for an SVM peer relation.

```

```
|cluster_peer
|link:#cluster_peer[cluster_peer]
a|Peer cluster details

|svm
|link:#svm[svm]
a|Local SVM permitted for peer relation. To create peer permissions for
all SVMs, specify the SVM name as "*".
```

```
|===
```

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "flexcache"
  ],
  "cluster_peer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster2",
    "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

```
}  
====  
  
== Error
```

Status: Default

ONTAP Error Response Codes

```
|===  
| Error codes | Description  
  
| 26345574  
| Failed to find the SVM or volume name with UUID.  
|===
```

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|error  
|link:#error[error]  
a|
```

```
|===
```

.Example error

```
[%collapsible%closed]
```

```
====
```

```
[source,json,subs=+macros]
```

```
{  
  "error": {  
    "arguments": [  
      {  
        "code": "string",  
        "message": "string"  
      }  
    ],  
    "code": "4",  
    "message": "entry doesn't exist",  
    "target": "uuid"
```

```

    }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#cluster_peer]
[.api-collapsible-fifth-title]

```

cluster_peer

Peer cluster details

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|
```

```
|uuid
```

```
|string
```

```
a|
```

```
|===
```

```
[#svm]
```

```
[.api-collapsible-fifth-title]
```

```
svm
```

Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "*".

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments
```

```

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID53fb5fbfef0dfa422bcf79b2f8eb7332]]
= Update SVM peer permissions

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/svm/peer-permissions/{cluster_peer.uuid}/{svm.uuid}`#

*Introduced In:* 9.6

Updates the SVM peer permissions.

== Related ONTAP commands

* `vserver peer permission modify`

== Example

Updates an SVM peer permission.

----
PATCH "/api/svm/peer-permissions/d3268a74-ee76-11e8-a9bb-
005056ac6dc9/8f467b93-f2f1-11e8-9027-005056ac81fc"
'{"applications":["flexcache"]}'
----
```

== Learn more

* xref:{relative_path}svm_peer-permissions_endpoint_overview.html[DOC
/svm/peer-permissions]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|cluster_peer.uuid

|string

|path

|True

a|Peer cluster UUID

* Introduced in: 9.7

|svm.uuid

|string

|path

|True

a|SVM UUID

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|applications

|array[string]

a|A list of applications for an SVM peer relation.

|===

.Example request

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "applications": [
    "snapmirror",
    "flexcache"
  ]
}
```

====

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|applications
|array[string]
a|A list of applications for an SVM peer relation.
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "applications": [
    "snapmirror",
    "flexcache"
  ]
}
====

== Error
```

Status: Default

ONTAP Error Response Codes

```
|===
| Error codes | Description

| 26345572
| \{field} is a required field.

| 26345574
| Failed to find the SVM or volume name with UUID.
|===
```

```
== Definitions
```

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#cluster_peer]
[.api-collapsible-fifth-title]
cluster_peer

Peer cluster details

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

Local SVM permitted for peer relation. To create peer permissions for all SVMs, specify the SVM name as "*".

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#svm_peer_permission]
[.api-collapsible-fifth-title]
svm_peer_permission
```

Manage SVM peer permissions.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|applications
|array[string]
a|A list of applications for an SVM peer relation.
```

```
|===
```

```
[#error_arguments]  
[.api-collapsible-fifth-title]  
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|code  
|string  
a|Argument code
```

```
|message  
|string  
a|Message argument
```

```
|===
```

```
[#error]  
[.api-collapsible-fifth-title]  
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|arguments  
|array[link:#error_arguments[error_arguments]]  
a|Message arguments
```

```
|code  
|string  
a|Error code
```

```
|message
```

```
|string
a|Error message

|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
:leveloffset: -1
```

```
= Manage SVM peer relationships
```

```
:leveloffset: +1
```

```
[[ID6f9a6884f816ef5c8d91abc41a0ec702]]
= SVM peers endpoint overview
```

```
== Managing SVM peers
```

The SVM peer commands allow you to create and manage SVM peering relationships.

```
=== SVM peer APIs
```

The following APIs are used to manage SVM peers:

```
* GET /api/svm/peers
* POST /api/svm/peers
* GET /api/svm/peers/{uuid}
* PATCH /api/svm/peers/{uuid}
* DELETE /api/svm/peers/{uuid}
```

```
[[IDb971aeaa2b78d3bca2f04612fff5ecc8]]
```

= Retrieve SVM peer relationships

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/svm/peers`#
```

Introduced In: 9.6

Retrieves the list of SVM peer relationships.

== Related ONTAP commands

* `vserver peer show`

== Examples

The following examples show how to retrieve a collection of SVM peer relationships based on a query.

. Retrieves a list of SVM peers of a specific local SVM

```
GET "/api/svm/peers/?svm.name=VS1"
```

. Retrieves a list of SVM peers of a specific cluster peer

```
GET "/api/svm/peers/?peer.cluster.name=cluster2"
```

== Learn more

* xref:{relative_path}svm_peers_endpoint_overview.html[DOC /svm/peers]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

|Description

|svm.uuid

|string

|query

|False

a|Filter by svm.uuid

|svm.name

|string

|query

|False

a|Filter by svm.name

|uuid

|string

|query

|False

a|Filter by uuid

|peer.cluster.name

|string

|query

|False

a|Filter by peer.cluster.name

|peer.cluster.uuid

|string

|query

|False

a|Filter by peer.cluster.uuid

|peer.svm.uuid

|string

|query

|False

a|Filter by peer.svm.uuid

|peer.svm.name

|string

|query


```
|False  
a|Filter by peer.svm.name
```

```
|applications  
|string  
|query  
|False  
a|Filter by applications
```

```
|name  
|string  
|query  
|False  
a|Filter by name
```

```
|state  
|string  
|query  
|False  
a|Filter by state
```

```
|fields  
|array[string]  
|query  
|False  
a|Specify the fields to return.
```

```
|max_records  
|integer  
|query  
|False  
a|Limit the number of records returned.
```

```
|return_records  
|boolean  
|query  
|False  
a|The default is true for GET calls. When set to false, only the number  
of records is returned.
```

```
* Default value: 1
```

```

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds.  ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|num_records
|integer
a|Number of records

|records
|array[link:#svm_peer[svm_peer]]

```

a|

|===

.Example response

[%collapsible%closed]

====

[source,json,subs=+macros]

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "applications": [
        "snapmirror",
        "lun_copy"
      ],
      "name": "string",
      "peer": {
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster2",
          "uuid": "ebe27c49-1adf-4496-8335-ab862aebbf2"
        },
        "svm": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          }
        }
      },
    }
  ],
}
```

```

        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "state": "peered",
    "svm": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "uuid": "string"
}
]
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error codes | Description
|
| 26345578
| Internal error. Unable to retrieve local or peer SVM name.
|===

```

[cols=3*,options=header]

```

|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

```

```

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]

```

```

[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:#href[href]
a|

|self
|link:#href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

SVM, applies only to SVM-scoped objects.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#peer]
[.api-collapsible-fifth-title]
```

peer

Details for a peer SVM object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.
```

|===

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

Local SVM details

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.
```

```
|uuid
|string
```


a|The unique identifier of the SVM.

|===

[#svm_peer]
[.api-collapsible-fifth-title]
svm_peer

An SVM peer relation object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|applications
|array[string]
a|A list of applications for an SVM peer relation.

|name
|string
a|A peer SVM alias name to avoid a name conflict on the local cluster.

|peer
|link:#peer[peer]
a|Details for a peer SVM object.

|state
|string
a|SVM peering state. To accept a pending SVM peer request, PATCH the state to "peered". To reject a pending SVM peer request, PATCH the state to "rejected". To suspend a peered SVM peer relation, PATCH the state to "suspended". To resume a suspended SVM peer relation, PATCH the state to "peered". The states "initiated", "pending", and "initializing" are system-generated and cannot be used for PATCH.

```
|svm
|link:#svm[svm]
a|Local SVM details
```

```
|uuid
|string
a|SVM peer relation UUID
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments


|code
|string
a|Error code


|message
|string
a|Error message


|target
|string
a|The target parameter that caused the error.


|===


//end collapsible .Definitions block
====


[[ID6564d8c5883c6b5cc11090270e02528c]]
= Create a new SVM peer relationship


[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-
block]#`/svm/peers`#


*Introduced In:* 9.6


Creates a new SVM peer relationship.


== Important notes


*** The create request accepts peer SVM name as input instead of peer SVM
UUID as the local cluster cannot validate peer SVM based on UUID.


*** The input parameter `name` refers to the local name of the peer SVM.
The `peer cluster name` parameter is optional for creating intracluster
SVM peer relationships.

```

== Required properties

- * `svm.name` or `svm.uuid` - SVM name or SVM UUID
- * `peer.svm.name` or `peer.svm.uuid` - Peer SVM name or Peer SVM UUID
- * `peer.cluster.name` or `peer.cluster.uuid` - Peer cluster name or peer cluster UUID
- * `applications` - Peering applications

== Related ONTAP commands

- * `vserver peer create`

== Example

Creates a new SVM peer relationship.

```
-----  
POST "/api/svm/peers" '{"svm":{"name":"vs1"},  
"peer.cluster.name":"cluster2", "peer.svm.name":"VS1",  
"applications":["snapmirror"]}'  
-----
```

== Learn more

- * [xref:{relative_path}svm_peers_endpoint_overview.html](#)[DOC /svm/peers]

== Parameters

```
[cols=5*,options=header]  
|===
```

```
|Name  
|Type  
|In  
|Required  
|Description
```

```
|return_timeout  
|integer  
|query  
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH,

or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

|return_records

|boolean

|query

|False

a|The default is false. If set to true, the records are returned.

* Default value:

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|applications

|array[string]

a|A list of applications for an SVM peer relation.

|name

|string

a|A peer SVM alias name to avoid a name conflict on the local cluster.

|peer

|link:#peer[peer]

a|Details for a peer SVM object.

|svm

|link:#svm[svm]

a|Local SVM details

```
|uuid
|string
a|SVM peer relation UUID
```

```
|===
```

```
.Example request
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "name": "string",
  "peer": {
    "cluster": {
      "name": "cluster2",
      "uuid": "ebe27c49-1adf-4496-8335-ab862aebbf2"
    },
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
=====
```

```
== Response
```

Status: 202, Accepted

```
[cols=3*,options=header]
|===
|Name
|Type
```

```

|Description

|applications
|array[string]
a|A list of applications for an SVM peer relation.

|name
|string
a|A peer SVM alias name to avoid a name conflict on the local cluster.

|peer
|link:#peer[peer]
a|Details for a peer SVM object.

|svm
|link:#svm[svm]
a|Local SVM details

|uuid
|string
a|SVM peer relation UUID

|===

```

```

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "name": "string",
  "peer": {
    "cluster": {
      "name": "cluster2",
      "uuid": "ebe27c49-1adf-4496-8335-ab862aebebf2"
    },
    "svm": {
      "name": "svm1",

```

```

        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
},
"svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "string"
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error codes | Description
|
| 13434889
| Internal error. Wait and retry.
|
| 26345575
| The specified peer cluster name and peer cluster UUID do not match.
|
| 26345579
| The specified field is invalid.
|
| 26345580
| SVM name or SVM UUID must be provided.
|===

```

== Definitions

```

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]

```



```

|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#peer]
[.api-collapsible-fifth-title]
peer

Details for a peer SVM object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

Local SVM details

```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.
```

```
|===

[#svm_peer]
[.api-collapsible-fifth-title]
svm_peer
```

An SVM peer relation object.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|applications
|array[string]
a|A list of applications for an SVM peer relation.
```

```
|name
|string
a|A peer SVM alias name to avoid a name conflict on the local cluster.
```

```
|peer
|link:#peer[peer]
a|Details for a peer SVM object.
```

```
|svm
|link:#svm[svm]
a|Local SVM details
```

```
|uuid
|string
a|SVM peer relation UUID
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|code
|string
a|Argument code
```

```
|message
|string
a|Message argument
```

```
|===
```

```
[#error]
[.api-collapsible-fifth-title]
error
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

```
a|Error message
```

```
|target
```

```
|string
```

```
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
```

```
====
```

```
[[ID810f934c54fd21992d4dffe1129193be]]
```

```
= Delete an SVM peer relationship
```

```
[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-  
block]#`/svm/peers/{uuid}`#
```

```
*Introduced In:* 9.6
```

```
Deletes the SVM peer relationship.
```

```
== Related ONTAP commands
```

```
* `vserver peer delete`
```

```
== Example
```

```
Deletes an SVM peer relationship.
```

```
----
```

```
DELETE "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
```

== Learn more

* xref:{relative_path}svm_peers_endpoint_overview.html[DOC /svm/peers]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|SVM peer relationship UUID

* Introduced in: 9.7

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

|===

== Response

Status: 202, Accepted

== Error

Status: Default

ONTAP Error Response Codes

|===

| Error codes | Description

| 26345578

| Internal error. Unable to retrieve local or peer SVM name.

|===

[cols=3*,options=header]

|===

|Name

|Type

|Description

|error

|link:#error[error]

a|

|===

.Example error

[%collapsible%closed]

====

[source,json,subs=+macros]

{

 "error": {

 "arguments": [

 {

 "code": "string",

 "message": "string"

 }

],

 "code": "4",

 "message": "entry doesn't exist",

 "target": "uuid"

```

    }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

```



```
|code
|string
a|Error code
```

```
|message
|string
a|Error message
```

```
|target
|string
a|The target parameter that caused the error.
```

```
|===
```

```
//end collapsible .Definitions block
====
```

```
[[IDc0207ca7fb06e421627d5c07fb966992]]
= Retrieve an SVM peer relationship instance
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-
block]#`/svm/peers/{uuid}`#
```

```
*Introduced In:* 9.6
```

Retrieves the SVM peer relationship instance.

== Related ONTAP commands

```
* `vserver peer show`
```

== Example

Retrieves the parameters of an SVM peer relationship.

```
----
GET "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
----
```

== Learn more

* xref:{relative_path}svm_peers_endpoint_overview.html[DOC /svm/peers]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|SVM peer relation UUID

* Introduced in: 9.7

|fields

|array[string]

|query

|False

a|Specify the fields to return.

|return_timeout

|integer

|query

|False

a|The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.

* Default value: 1

* Max value: 120

* Min value: 0

|===

== Response

Status: 200, Ok

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|applications
|array[string]
a|A list of applications for an SVM peer relation.

|name
|string
a|A peer SVM alias name to avoid a name conflict on the local cluster.

|peer
|link:#peer[peer]
a|Details for a peer SVM object.

|state
|string
a|SVM peering state. To accept a pending SVM peer request, PATCH the state
to "peered". To reject a pending SVM peer request, PATCH the state to
"rejected". To suspend a peered SVM peer relation, PATCH the state to
"suspended". To resume a suspended SVM peer relation, PATCH the state to
"peered". The states "initiated", "pending", and "initializing" are
system-generated and cannot be used for PATCH.

|svm
|link:#svm[svm]
a|Local SVM details

|uuid
|string
```

a|SVM peer relation UUID

|===

.Example response

[%collapsible%closed]

=====

[source,json,subs=+macros]

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "name": "string",
  "peer": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster2",
      "uuid": "ebe27c49-1adf-4496-8335-ab862aebbf2"
    },
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "peered",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```

```

    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "string"
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|====
| Error codes | Description
|
| 26345578
| Internal error. Unable to retrieve local or peer SVM name.
|====

```

[cols=3*,options=header]

```

|====
|Name
|Type
|Description

```

```

|error
|link:#error[error]
a|

```

```

|====

```

.Example error

[%collapsible%closed]

```

====

```

[source,json,subs=+macros]

```

{
  "error": {
    "arguments": [
      {
        "code": "string",

```

```

        "message": "string"
    }
],
"code": "4",
"message": "entry doesn't exist",
"target": "uuid"
}
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]

```

```

a|

|===

[#cluster]
[.api-collapsible-fifth-title]
cluster

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#peer]
[.api-collapsible-fifth-title]
peer

Details for a peer SVM object.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|cluster
|link:#cluster[cluster]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

Local SVM details

[cols=3*,options=header]
|===

```



```

|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|The name of the SVM.

|uuid
|string
a|The unique identifier of the SVM.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID9ec8fdb9a94883921946136eb1f42b00]]
= Update an SVM peer relationship

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/svm/peers/{uuid}`#

*Introduced In:* 9.6

Updates the SVM peer relationship.

== Related ONTAP commands

```

* `vserver peer modify`

== Examples

The following examples show how to update an SVM peer relationship. The input parameter 'name' refers to the local name of the peer SVM.

. Accepts an SVM peer relationship

```
PATCH "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
'{"state":"peered"}'
```

. Updates the local name of an SVM peer relationship

```
PATCH "/api/svm/peers/d3268a74-ee76-11e8-a9bb-005056ac6dc9"
'{"name":"vs2"}'
```

== Learn more

* xref:{relative_path}svm_peers_endpoint_overview.html[DOC /svm/peers]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|SVM peer relationship UUID

* Introduced in: 9.7

|return_timeout

```
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When doing a POST, PATCH, or DELETE operation on a single record, the
default is 0 seconds. This means that if an asynchronous operation is
started, the server immediately returns HTTP code 202 (Accepted) along
with a link to the job. If a non-zero value is specified for POST, PATCH,
or DELETE operations, ONTAP waits that length of time to see if the job
completes so it can return something other than 202.
```

```
* Default value: 1
* Max value: 120
* Min value: 0
```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|applications
```

```
|array[string]
```

```
a|A list of applications for an SVM peer relation.
```

```
|state
```

```
|string
```

```
a|SVM peering state. To accept a pending SVM peer request, PATCH the state
to "peered". To reject a pending SVM peer request, PATCH the state to
"rejected". To suspend a peered SVM peer relation, PATCH the state to
"suspended". To resume a suspended SVM peer relation, PATCH the state to
"peered". The states "initiated", "pending", and "initializing" are
system-generated and cannot be used for PATCH.
```

```
|uuid
```

```
|string
```

```
a|SVM peer relation UUID
```

```

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "applications": [
    "snapmirror",
    "lun_copy"
  ],
  "state": "peered",
  "uuid": "string"
}
====

== Response

```

Status: 202, Accepted

```

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error codes | Description

| 13434889
| Internal error. Wait and retry.

| 26345575
| The specified peer cluster name and peer cluster UUID do not match.

| 26345576
| Given peer state is invalid.

| 26345577
| One of the following is required: applications, state, or name.

| 26345578
| Internal error. Unable to retrieve local or peer SVM name.

```

```
| 26345579
| The specified field is invalid.

| 26345581
| Peer cluster name could not be retrieved or validated.
|===
```

== Definitions

```
[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
```

```
[#href]
[.api-collapsible-fifth-title]
href
```

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|href
|string
a|
```

```
|===
```

```
[#_links]
[.api-collapsible-fifth-title]
_links
[#cluster]
[.api-collapsible-fifth-title]
cluster
```

```
[cols=3*,options=header]
```

```
|===
|Name
|Type
|Description
```

```
|name
|string
```

```

a|

|uuid
|string
a|

|===

[#svm]
[.api-collapsible-fifth-title]
svm

SVM, applies only to SVM-scoped objects.

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|name
|string
a|The name of the SVM.

```

```

|uuid
|string
a|The unique identifier of the SVM.

```

```

|===

```

```

[#peer]
[.api-collapsible-fifth-title]
peer

```

Details for a peer SVM object.

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```
|cluster
|link:#cluster[cluster]
a|

|svm
|link:#svm[svm]
a|SVM, applies only to SVM-scoped objects.
```

```
|===
```

```
[#svm]
[.api-collapsible-fifth-title]
svm
```

Local SVM details

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|The name of the SVM.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|===
```

```
[#svm_peer]
[.api-collapsible-fifth-title]
svm_peer
```

An SVM peer relation object.

```
[cols=3*,options=header]
```



```

|===
|Name
|Type
|Description

|applications
|array[string]
a|A list of applications for an SVM peer relation.


|state
|string
a|SVM peering state. To accept a pending SVM peer request, PATCH the state
to "peered". To reject a pending SVM peer request, PATCH the state to
"rejected". To suspend a peered SVM peer relation, PATCH the state to
"suspended". To resume a suspended SVM peer relation, PATCH the state to
"peered". The states "initiated", "pending", and "initializing" are
system-generated and cannot be used for PATCH.


|uuid
|string
a|SVM peer relation UUID


|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code


|message
|string
a|Message argument

```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

= Manage SVMs

```

```
:leveloffset: +1
```

```
[[IDa125addf1709a0dd76dbba162e362f21]]  
= SVM svms endpoint overview
```

== Managing SVMs

Cluster administrators can manage any SVM bound to the cluster. In addition, SVMs can also be managed by their SVM administrators. The SVM administrator manages the SVM resources, such as volumes, protocols and services, depending on the capabilities assigned by the cluster administrator. SVM administrators cannot create, modify, or delete SVMs. The cluster administrator manages SVM create, modify, or delete operations.

While configuring CIFS, you must also configure IP interfaces and DNS. No other protocol configuration is allowed when configuring NVMe. NFS, FCP, CIFS and iSCSI protocols can be configured together.

SVM administrators might have all or some of the following administration capabilities:

- . Data access protocol configuration

Configures data access protocols, such as NFS, CIFS, iSCSI, and Fibre Channel (FC) protocol (Fibre Channel over Ethernet included).

- . Services configuration

Configures services such as LDAP, NIS, and DNS.

- . Monitoring SVM

Monitors jobs, network connections, network interfaces, and SVM health.

- . Updating the TLS certificate for this SVM.

```
[[ID0aab82f414295ad97374a4a9c7d393c9]]  
= Retrieve SVMs and SVM properties
```

```
[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-  
block]#`/svm/svms`#
```

Introduced In: 9.6

Retrieves a list of SVMs and individual SVM properties. This includes

protocol configurations such as CIFS and NFS, export policies, name service configurations, and network services.

== Important notes

* The SVM object includes a large set of fields and can be expensive to retrieve. Use this API to list the collection of SVMs, and to retrieve only the full details of individual SVMs as needed.

* It is not recommended to create or delete more than five SVMs in parallel.

* REST APIs only expose a data SVM as an SVM.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `snapmirror.+*+`

== Related ONTAP commands

* `vserver show`

== Examples

. Retrieves a list of SVMs in the cluster sorted by name

```
GET "/api/svm/svms?order_by=name"
```

. Retrieves a list of SVMs in the cluster that have the NFS protocol enabled

```
GET "/api/svm/svms?nfs.enabled=true"
```

. Retrieves a list of SVMs in the cluster that have the CIFS protocol enabled

```
GET "/api/svm/svms?cifs.enabled=true"
```

. Retrieves a list of SVMs in the cluster that have the S3 protocol enabled

GET "/api/svm/svms?s3.enabled=true"

== Learn more

* xref:{relative_path}svm_svms_endpoint_overview.html[DOC /svm/svms]

== Parameters

[cols=5*,options=header]

|==

|Name

|Type

|In

|Required

|Description

|nis.enabled

|boolean

|query

|False

a|Filter by nis.enabled

|nis.servers

|string

|query

|False

a|Filter by nis.servers

|nis.domain

|string

|query

|False

a|Filter by nis.domain

|nvme.enabled

|boolean

```
|query
|False
a|Filter by nvme.enabled
```

```
|language
|string
|query
|False
a|Filter by language
```

```
|nfs.enabled
|boolean
|query
|False
a|Filter by nfs.enabled
```

```
|comment
|string
|query
|False
a|Filter by comment
```

```
|aggregates.name
|string
|query
|False
a|Filter by aggregates.name
```

```
|aggregates.uuid
|string
|query
|False
a|Filter by aggregates.uuid
```

```
|subtype
|string
|query
|False
a|Filter by subtype
```

```
|dns.servers
|string
|query
|False
a|Filter by dns.servers
```

```
|dns.domains
|string
|query
|False
a|Filter by dns.domains
```

```
|fcp.enabled
|boolean
|query
|False
a|Filter by fcp.enabled
```

```
|iscsi.enabled
|boolean
|query
|False
a|Filter by iscsi.enabled
```

```
|name
|string
|query
|False
a|Filter by name
```

```
|ipSPACE.uuid
|string
|query
|False
a|Filter by ipSPACE.uuid
```

```
|ipSPACE.name
|string
|query
|False
a|Filter by ipSPACE.name
```

```
|ldap.base_dn
|string
|query
|False
a|Filter by ldap.base_dn
```

```
|ldap.servers
|string
|query
|False
a|Filter by ldap.servers
```

```
|ldap.enabled
|boolean
|query
|False
a|Filter by ldap.enabled
```

```
|ldap.bind_dn
|string
|query
|False
a|Filter by ldap.bind_dn
```

```
|ldap.ad_domain
|string
|query
|False
a|Filter by ldap.ad_domain
```

```
|uuid
|string
|query
|False
a|Filter by uuid
```

```
|cifs.name
|string
|query
```



```
|False  
a|Filter by cifs.name
```

```
|cifs.ad_domain.fqdn  
|string  
|query  
|False  
a|Filter by cifs.ad_domain.fqdn
```

```
|cifs.ad_domain.organizational_unit  
|string  
|query  
|False  
a|Filter by cifs.ad_domain.organizational_unit
```

```
|cifs.enabled  
|boolean  
|query  
|False  
a|Filter by cifs.enabled
```

```
|s3.name  
|string  
|query  
|False  
a|Filter by s3.name
```

* Introduced in: 9.7

```
|s3.enabled  
|boolean  
|query  
|False  
a|Filter by s3.enabled
```

* Introduced in: 9.7

```
|nsswitch.netgroup  
|string  
|query  
|False
```

```
a|Filter by nsswitch.netgroup
```

```
|nsswitch.group
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by nsswitch.group
```

```
|nsswitch.hosts
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by nsswitch.hosts
```

```
|nsswitch.namemap
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by nsswitch.namemap
```

```
|nsswitch.passwd
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by nsswitch.passwd
```

```
|fc_interfaces.name
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by fc_interfaces.name
```

```
* Introduced in: 9.7
```

```
|fc_interfaces.data_protocol
```

```
|string
```

```
|query
```

```
|False
```

```
a|Filter by fc_interfaces.data_protocol
```

```
* Introduced in: 9.7
```

```
|snapshot_policy.uuid
|string
|query
|False
a|Filter by snapshot_policy.uuid
```

```
|snapshot_policy.name
|string
|query
|False
a|Filter by snapshot_policy.name
```

```
|state
|string
|query
|False
a|Filter by state
```

```
|certificate.uuid
|string
|query
|False
a|Filter by certificate.uuid
```

* Introduced in: 9.7

```
|fields
|array[string]
|query
|False
a|Specify the fields to return.
```

```
|max_records
|integer
|query
|False
a|Limit the number of records returned.
```

```
|return_records
```

```

|boolean
|query
|False
a|The default is true for GET calls. When set to false, only the number
of records is returned.

* Default value: 1

|return_timeout
|integer
|query
|False
a|The number of seconds to allow the call to execute before returning.
When iterating over a collection, the default is 15 seconds. ONTAP
returns earlier if either max records or the end of the collection is
reached.

* Default value: 1
* Max value: 120
* Min value: 0

|order_by
|array[string]
|query
|False
a|Order results by specified fields and optional [asc|desc] direction.
Default direction is 'asc' for ascending.

|===

== Response

```

Status: 200, Ok

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

```

```
|num_records
|integer
a|Number of records
```

```
|records
|array[link:#svm[svm]]
a|
```

```
|===
```

```
.Example response
[%collapsible%closed]
```

```
=====
```

```
[source,json,subs=+macros]
```

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "certificate": {
        "_links": {
          "self": {
```

```

        "href": "/api/resourcelink"
    },
    },
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"cifs": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    },
    "ad_domain": {
        "fqdn": "example.com",
        "organizational_unit": "string"
    },
    },
    "name": "CIFS1"
},
"comment": "string",
"dns": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    },
    "domains": [
        "example.com",
        "example2.example3.com"
    ],
    "servers": [
        "10.224.65.20",
        "2001:db08:a0b:12f0::1"
    ]
},
"fc_interfaces": [
    {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        },
        "data_protocol": "string",
        "name": "lif1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
],

```

```

"fcg": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"ip_interfaces": [
  {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "ip": {
      "address": "10.10.10.7"
    },
    "name": "lif1",
    "services": [
      "data_nfs"
    ],
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"ipspace": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "exchange",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"iscsi": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"language": "c.utf_8",
"ldap": {
  "ad_domain": "string",
  "base_dn": "string",
  "bind_dn": "string",
  "servers": [

```

```

        "string"
    ],
    },
    "name": "svm1",
    "nfs": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        }
    },
    },
    "nis": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        }
    },
    "domain": "string",
    "servers": [
        "string"
    ],
    },
    "nsswitch": {
        "group": [
            "string"
        ],
        "hosts": [
            "string"
        ],
        "namemap": [
            "string"
        ],
        "netgroup": [
            "string"
        ],
        "passwd": [
            "string"
        ],
    },
    },
    "nvme": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        }
    },
    },
    },

```



```

"s3": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "s3-server-1"
},
"snapmirror": {
  "protected_volumes_count": 0
},
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"state": "running",
"subtype": "string",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7",
"volume_efficiency_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
]
}
=====

```

== Error

Status: Default, Error

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href
```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|next
|link:href[href]
a|

|self
|link:href[href]
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:href[href]

```

```

a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#certificate]
[.api-collapsible-fifth-title]
certificate

Certificate for incoming TLS connection requests.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links

```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|Certificate name
```

```
|uuid
```

```
|string
```

```
a|Certificate UUID
```

```
|===
```

```
[#ad_domain]
```

```
[.api-collapsible-fifth-title]
```

```
ad_domain
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|fqdn
```

```
|string
```

```
a|The fully qualified domain name of the Windows Active Directory to which  
this CIFS server belongs. A CIFS server appears as a member of Windows  
server object in the Active Directory store.
```

```
|organizational_unit
```

```
|string
```

```
a|Specifies the organizational unit within the Active Directory domain to  
associate with the CIFS server.
```

```
|===
```

```
[#cifs]
```

```
[.api-collapsible-fifth-title]
```

```
cifs
```

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ad_domain
|link:#ad_domain[ad_domain]
a|

|enabled
|boolean
a|Specifies whether or not the CIFS service is administratively enabled.

|name
|string
a|The NetBIOS name of the CIFS server.

|===

[#dns]
[.api-collapsible-fifth-title]
dns

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|domains
|array[string]
a|A list of DNS domains.
Domain names have the following requirements:

* The name must contain only the following characters: A through Z,
a through z, 0 through 9, ".", "-", or "_".

```

- * The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.
- * The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.
- * The top level domain must contain only the following characters: A through Z, a through z.
- * The system reserves the following names: "all", "local", and "localhost".

```
|servers
|array[string]
a|The list of IP addresses of the DNS servers. Addresses can be either
IPv4 or IPv6 addresses.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#fc_port_reference]
[.api-collapsible-fifth-title]
fc_port_reference
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|The name of the FC port.
```

```
|node
```

```
|link:#node[node]
```

```
a|The node on which the FC port is located.
```

```
|uuid
```

```
|string
```

```
a|The unique identifier of the FC port.
```

```
|===
```

```
[#location]
```

```
[.api-collapsible-fifth-title]
```

```
location
```

The location of the Fibre Channel interface is defined by the location of its port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|port
```

```
|link:#fc_port_reference[fc_port_reference]
```

```
a|An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.
```



```
|===
```

```
[#fc_interface_svm]  
[.api-collapsible-fifth-title]  
fc_interface_svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|data_protocol  
|string  
a|The data protocol for which the Fibre Channel interface is configured.
```

```
|name  
|string  
a|The name of the Fibre Channel interface.
```

```
|uuid  
|string  
a|The unique identifier of the Fibre Channel interface.
```

```
|===
```

```
[#fcp]  
[.api-collapsible-fifth-title]  
fcp
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name  
|Type  
|Description
```

```

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|Enable Fiber Channel Protocol (FCP)? Setting to true creates a service
if not already created.

|===

[#ip]
[.api-collapsible-fifth-title]
ip

IP information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|===

[#broadcast_domain]
[.api-collapsible-fifth-title]
broadcast_domain

Broadcast domain UUID along with a readable name.

[cols=3*,options=header]
|===
|Name
|Type
|Description

```

```

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the broadcast domain, scoped to its IPspace

|uuid
|string
a|Broadcast domain UUID

|===

[#home_node]
[.api-collapsible-fifth-title]
home_node

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|

|uuid
|string
a|

|===

[#location]
[.api-collapsible-fifth-title]
location

Home_node is optional.

```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description

|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name.
```

```
|home_node
|link:#home_node[home_node]
a|
```

```
|===
```

```
[#ip_interface_svm]
[.api-collapsible-fifth-title]
ip_interface_svm
```

Interface parameters. Name and home_node are optional.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|ip
|link:#ip[ip]
a|IP information
```

```
|name
|string
a|The name of the interface (optional).
```

```

|services
|array[string]
a|The services associated with the interface.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

Either the UUID or name may be supplied on input.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|IPspace name

|uuid
|string
a|IPspace UUID

|===

[#iscsi]
[.api-collapsible-fifth-title]
iscsi

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|Enable iSCSI? Setting to true creates a service if not already created.

|===

[#ldap]
[.api-collapsible-fifth-title]
ldap

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_domain
|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST.

|base_dn
|string
a|Specifies the default base DN for all searches.

|bind_dn
|string
a|Specifies the user that binds to the LDAP servers. SVM API supports
anonymous binding. For Simple and SASL LDAP binding, use the LDAP API
endpoint.

```

```

|enabled
|boolean
a|Enable LDAP? Setting to true creates a configuration if not already
created.

|servers
|array[string]
a|

|===

[#nfs]
[.api-collapsible-fifth-title]
nfs

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|Enable NFS? Setting to true creates a service if not already created.

|===

[#nis]
[.api-collapsible-fifth-title]
nis

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]

```

```

a|

|domain
|string
a|The NIS domain to which this configuration belongs.

|enabled
|boolean
a|Enable NIS? Setting to true creates a configuration if not already
created.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|===

[#nsswitch]
[.api-collapsible-fifth-title]
nsswitch

Name service switch configuration

[cols=3*,options=header]
|===
|Name
|Type
|Description

|group
|array[string]
a|Group sources

|hosts
|array[string]
a|Host sources

|namemap
|array[string]

```



```
a|NameMap sources
```

```
|netgroup  
|array[string]  
a|NetGroup sources
```

```
|passwd  
|array[string]  
a|Password sources
```

```
|===
```

```
[#nvme]  
[.api-collapsible-fifth-title]  
nvme
```

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|_links  
|link:#_links[_links]  
a|
```

```
|enabled  
|boolean  
a|Enable NVMe? Setting to true creates a service if not already created.
```

```
|===
```

```
[#ip_info]  
[.api-collapsible-fifth-title]  
ip_info
```

```
IP information
```

```
[cols=3*,options=header]  
|===
```

```

|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|family
|string
a|IPv4 or IPv6

|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.

|===

[#network_route_for_svm]
[.api-collapsible-fifth-title]
network_route_for_svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|destination
|link:#ip_info[ip_info]
a|IP information

|gateway
|string
a|The IP address of the gateway router leading to the destination.

|===

```

```
[#certificate]
[.api-collapsible-fifth-title]
certificate
```

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID
```

```
|===
```

```
[#s3]
[.api-collapsible-fifth-title]
s3
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|certificate
|link:#certificate[certificate]
```

a|Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

|enabled

|boolean

a|Specifies whether or not to enable S3. Setting this value to true creates a service if one is not yet created.

|is_http_enabled

|boolean

a|Specifies whether HTTP is enabled on the S3 server. By default, HTTP is disabled on the S3 server.

|is_https_enabled

|boolean

a|Specifies whether HTTPS is enabled on the S3 server. By default, HTTPS is enabled on the S3 server.

|name

|string

a|Specifies the name of the S3 server. A server name length can range from 1 to 15 characters and can only contain the following combination of characters 0-9, A-Z, a-z, ".", and "-".

|port

|integer

a|Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80.

|secure_port

|integer

a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443.

|===

[#snapmirror]

[.api-collapsible-fifth-title]

snapmirror

Specifies attributes for SVM DR protection.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|is_protected
```

```
|boolean
```

a|Specifies whether the SVM is a SnapMirror source SVM, using SnapMirror to protect its data.

```
|protected_volumes_count
```

```
|integer
```

a|Specifies the number of SVM DR protected volumes in the SVM.

```
|===
```

```
[#snapshot_policy]
```

```
[.api-collapsible-fifth-title]
```

snapshot_policy

This is a reference to the Snapshot copy policy.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

a|

```
|name
```

```
|string
```

a|

```
|uuid
```

```
|string
```

```

a|

|===

[#volume_efficiency_policy]
[.api-collapsible-fifth-title]
volume_efficiency_policy

This is a reference to the volume efficiency policy.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the volume efficiency policy


|uuid
|string
a|Unique identifier of the volume efficiency policy.


|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]

```

```

a|

|aggregates
|array[link:#aggregates[aggregates]]
a|List of allowed aggregates for SVM volumes. An administrator is allowed
to create volumes on these aggregates.

|aggregates_delegated
|boolean
a|This property is true when the administrator has delegated the
aggregates for the SVM volumes.

|certificate
|link:#certificate[certificate]
a|Certificate for incoming TLS connection requests.

|cifs
|link:#cifs[cifs]
a|

|comment
|string
a|Comment

|dns
|link:#dns[dns]
a|

|fc_interfaces
|array[link:#fc_interface_svm[fc_interface_svm]]
a|FC Interface for the SVM

|fcp
|link:#fcf[fcf]
a|

|ip_interfaces
|array[link:#ip_interface_svm[ip_interface_svm]]
a|IP interfaces for the SVM

|ipspace

```

|link:#ipspace[ipspace]

a|Either the UUID or name may be supplied on input.

|iscsi

|link:#iscsi[iscsi]

a|

|language

|string

a|Default volume language code. UTF-8 encoded languages are valid in POST or PATCH. Non UTF-8 language encodings are for backward compatibility and are not valid input for POST and PATCH requests.

|ldap

|link:#ldap[ldap]

a|

|name

|string

a|The name of the SVM.

|nfs

|link:#nfs[nfs]

a|

|nis

|link:#nis[nis]

a|

|nsswitch

|link:#nsswitch[nsswitch]

a|Name service switch configuration

|nvme

|link:#nvme[nvme]

a|

|s3

|link:#s3[s3]

a|

|snapmirror

|link:#snapmirror[snapmirror]

a|Specifies attributes for SVM DR protection.

|snapshot_policy

|link:#snapshot_policy[snapshot_policy]

a|This is a reference to the Snapshot copy policy.

|state

|string

a|SVM State

|subtype

|string

a|SVM subtype. The SVM subtype sync_destination is created automatically when an SVM of subtype sync_source is created on the source MetroCluster cluster. A POST request with sync_destination as SVM subtype is invalid.

|uuid

|string

a|The unique identifier of the SVM.

|volume_efficiency_policy

|link:#volume_efficiency_policy[volume_efficiency_policy]

a|This is a reference to the volume efficiency policy.

|===

[#error_arguments]

[.api-collapsible-fifth-title]

error_arguments

[cols=3*,options=header]

|===

|Name

|Type

|Description

|code

|string

a|Argument code

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
=====

```

```
[[IDfcabe8fd9c24349fa12158a505050aae]]
```

= Create and provision an SVM

```
[.api-doc-operation .api-doc-operation-post]#POST# [.api-doc-code-block]#`/svm/svms`#
```

Introduced In: 9.6

Creates and provisions an SVM. If no IPspace is provided, then the SVM is created on the `Default` IPspace.

* The number of parallel SVMs that can be created must not be greater than five.

* If a sixth SVM POST request is issued, the following error message is generated: "Maximum allowed SVM jobs exceeded. Wait for the existing SVM jobs to complete and try again."

== Required properties

* `name` - Name of the SVM to be created.

== Recommended optional properties

* `ipspace.name` or `ipspace.uuid` - IPspace of the SVM

* `ip_interfaces` - If provided, the following fields are required:

* `ip_interfaces.name` - Name of the interface

* `ip_interfaces.ip.address` - IP address

* `ip_interfaces.ip.netmask` - Netmask length or IP address

* `ip_interfaces.location.broadcast_domain.uuid` or

`ip_interfaces.location.broadcast_domain.name` - Broadcast domain name or UUID belonging to the same IPspace of the SVM.

* `routes` - If provided, the following field is required:

*** `routes.gateway` - Gateway IP address

* `cifs` - If provided, interfaces, routes and DNS must be provided. The following fields are also required:

*** `cifs.name` - Name of the CIFS server to be created for the SVM.

*** `cifs.ad_domain.fqdn` - Fully qualified domain name

*** `cifs.ad_domain.user` - Administrator username

*** `cifs.ad_domain.password` - User password

* `ldap` - If provided, the following fields are required:

```

*** `ldap.servers` or `ldap.ad_domain` - LDAP server list or Active
Directory domain

*** `ldap.bind_dn` - Bind DN

*** `ldap.base_dn` - Base DN

* `nis` - If provided, the following fields are required:

*** `nis.servers` - NIS servers

*** `nis.domain` - NIS domain

* `dns` - If provided, the following fields are required:

*** `dns.servers` - Name servers

*** `dns.domains` - Domains

* `fc_interfaces` - If provided, the following fields are required:

*** `fc_interfaces.name` - Fibre Channel interface name

*** `fc_interfaces.data_protocol` - Fibre Channel interface data protocol

*** `fc_interfaces.location.port.uuid` or
`fc_interfaces.location.port.name` and
`fc_interfaces.location.port.node.name` - Either port UUID or port name
and node name together must be provided.

* `s3` - If provided, the following field should also be specified:

*** `s3.name` - Name of the S3 server. If ``s3.name` is not specified
while ``s3.enabled` is set to 'true', the S3 server will be created with
the default name '+++<svm.name>+++_S3Server'.+++</svm.name>+++

```

== Default property values

If not specified in POST, the following default property values are assigned:

```

* `language` - _C.UTF-8_
* `ipspace.name` - _Default_
* `snapshot_policy.name` - _Default_
* `subtype` - _Default_ ( _sync-source_ if MetroCluster configuration )

```

== Related ONTAP commands

```
* `vserver create`
* `vserver add-aggregates`
* `network interface create`
* `network route create`
* `vserver services name-service dns create`
* `vserver nfs create`
* `vserver services name-service ldap client create`
* `vserver cifs create`
* `vserver services name-service nis-domain create`
* `vserver iscsi create`
* `vserver nvme create`
* `vserver fcp create`
* `vserver services name-service ns-switch create`
* `vserver object-store-server create`
```

== Examples

. Creates an SVM with default "snapshot_policy"

```
-----
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"}}'
-----
```

. Creates an SVM and configures NFS, ISCSI and FCP

```
-----
POST "/api/svm/svms" '{"name":"testVs", "nfs":{"enabled":"true"},
"fcp":{"enabled":"true"}, "iscsi":{"enabled":"true"}}'
-----
```

. Creates an SVM and configures NVMe

```
-----
POST "/api/svm/svms" '{"name":"testVs", "nvme":{"enabled":"true"}}'
-----
```

. Creates an SVM and configures LDAP

```
-----
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"},
"ldap":{"servers":["10.140.101.1","10.140.101.2"], "ad_domain":"abc.com",
"base_dn":"dc=netapp,dc=com", "bind_dn":"dc=netapp,dc=com"}}'
-----
```

. Creates an SVM and configures NIS

```
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"}, "nis":{"enabled":"true",
"domain":"def.com","servers":["10.224.223.130", "10.224.223.131"]}]}'
```

. Creates an SVM and configures DNS

```
POST "/api/svm/svms" '{"name":"testVs",
"snapshot_policy":{"name":"default"},
"dns":{"domains":["abc.com","def.com"], "servers":["10.224.223.130",
"10.224.223.131"]}]}'
```

. Creates an SVM and configures a LIF

```
POST "/api/svm/svms" '{"name":"testVs", "ip_interfaces":
[{"name":"lif1", "ip":{"address":"10.10.10.7", "netmask":
"255.255.255.0"}, "location":{"broadcast_domain":{"name":"bd1"},
"home_node":{"name":"node1"}}, "service_policy": "default-management"]}]'
```

. Creates an SVM and configures a LIF with IPV6 address

```
POST "/api/svm/svms" '{"name":"testVs", "ip_interfaces":
[{"name":"lif2", "ip":{"address":"fd22:8b1e:b255:202:2a0:98ff:fe01:7d5b",
"netmask":"24"}, "location":{"broadcast_domain":{"name":"bd1"},
"home_node":{"name":"node1"}}, "service_policy": "default-management"]}]'
```

. Creates an SVM and configures CIFS

```
POST "/api/svm/svms" '{"name":"testVs", "cifs":{"name":"CIFDOC",
"ad_domain":{"fqdn":"abc.def.com", "organizational_unit":"CN=Computers",
"user":"cif_admin", "password":"abc123"}},
"ip_interfaces":[{"name":"lif1", "ip":{"address":"10.10.10.7", "netmask":
"255.255.255.0"}, "location":{"broadcast_domain":{"name":"bd1"},
"home_node":{"name":"node1"}}, "service_policy": "default-
management"}], "routes": [{"destination": {"address": "0.0.0.0", "netmask":
```

```
"0"}, "gateway": "10.10.10.7"}], "dns":{"domains":["abc.def.com",
"def.com"], "servers":["10.224.223.130", "10.224.223.131"]}}'
----

. Creates an SVM and configures an S3 server

----

POST "/api/svm/svms" '{"name":"svm5", "s3":{"name":"s3-server-1",
"enabled":true}}'
----

== Learn more

* xref:{relative_path}svm_svms_endpoint_overview.html[DOC /svm/svms]
```

== Parameters

```
[cols=5*,options=header]
|===
```

```
|Name
|Type
|In
|Required
|Description
```

```
|return_records
|boolean
|query
|False
```

a|The default is false. If set to true, the records are returned.

* Default value:

```
|return_timeout
|integer
|query
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

```
* Default value: 1
* Max value: 120
* Min value: 0
```

```
|===
```

```
== Request Body
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|aggregates
```

```
|array[link:#aggregates[aggregates]]
```

```
a|List of allowed aggregates for SVM volumes. An administrator is allowed
to create volumes on these aggregates.
```

```
|aggregates_delegated
```

```
|boolean
```

```
a|This property is true when the administrator has delegated the
aggregates for the SVM volumes.
```

```
|cifs
```

```
|link:#cifs[cifs]
```

```
a|
```

```
|comment
```

```
|string
```

```
a|Comment
```

```
|dns
```

```
|link:#dns[dns]
```

```
a|
```

```
|fc_interfaces
```

```
|array[link:#fc_interface_svm[fc_interface_svm]]
```

```
a|FC Interface for the SVM
```



```
| fcp
| link: #fcp[fcp]
a|

| ip_interfaces
| array[link: #ip_interface_svm[ip_interface_svm]]
a| IP interfaces for the SVM

| ipspace
| link: #ipspace[ipspace]
a| Either the UUID or name may be supplied on input.

| iscsi
| link: #iscsi[iscsi]
a|

| language
| string
a| Default volume language code. UTF-8 encoded languages are valid in POST
or PATCH. Non UTF-8 language encodings are for backward compatibility and
are not valid input for POST and PATCH requests.

| ldap
| link: #ldap[ldap]
a|

| name
| string
a| The name of the SVM.

| nfs
| link: #nfs[nfs]
a|

| nis
| link: #nis[nis]
a|

| nsswitch
| link: #nsswitch[nsswitch]
a| Name service switch configuration
```

```

|nvme
|link:#nvme[nvme]
a|

|routes
|array[link:#network_route_for_svm[network_route_for_svm]]
a|Optional array of routes for the SVM

|s3
|link:#s3[s3]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SVM DR protection.

|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the Snapshot copy policy.

|subtype
|string
a|SVM subtype. The SVM subtype sync_destination is created automatically
when an SVM of subtype sync_source is created on the source MetroCluster
cluster. A POST request with sync_destination as SVM subtype is invalid.

|uuid
|string
a|The unique identifier of the SVM.

|volume_efficiency_policy
|link:#volume_efficiency_policy[volume_efficiency_policy]
a|This is a reference to the volume efficiency policy.

|===

.Example request
[%collapsible%closed]
====
[source,json,subs=+macros]

```

```

{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "cifs": {
    "ad_domain": {
      "fqdn": "example.com",
      "organizational_unit": "string",
      "password": "string",
      "user": "string"
    },
    "name": "CIFS1"
  },
  "comment": "string",
  "dns": {
    "domains": [
      "example.com",
      "example2.example3.com"
    ],
    "servers": [
      "10.224.65.20",
      "2001:db08:a0b:12f0::1"
    ]
  },
  "fc_interfaces": [
    {
      "data_protocol": "string",
      "location": {
        "port": {
          "name": "0a",
          "node": {
            "name": "node1"
          },
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      },
      "name": "lif1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "ip_interfaces": [
    {
      "ip": {

```

```

    "address": "10.10.10.7",
    "netmask": "24"
  },
  "location": {
    "broadcast_domain": {
      "name": "bd1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "name": "lif1",
  "service_policy": "string",
  "services": [
    "data_nfs"
  ],
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"ipspace": {
  "name": "exchange",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"language": "c.utf_8",
"ldap": {
  "ad_domain": "string",
  "base_dn": "string",
  "bind_dn": "string",
  "servers": [
    "string"
  ]
},
"name": "svml",
"nis": {
  "domain": "string",
  "servers": [
    "string"
  ]
},
"nsswitch": {
  "group": [
    "string"
  ],
  "hosts": [

```

```

        "string"
    ],
    "namemap": [
        "string"
    ],
    "netgroup": [
        "string"
    ],
    "passwd": [
        "string"
    ]
},
"routes": [
    {
        "destination": {
            "address": "10.10.10.7",
            "netmask": "24"
        },
        "gateway": "10.1.1.1"
    }
],
"s3": {
    "certificate": {
        "name": "cert1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "name": "s3-server-1"
},
"snapmirror": {
    "protected_volumes_count": 0
},
"snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"subtype": "string",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7",
"volume_efficiency_policy": {
    "name": "default",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
}
}
====

```

== Response

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error codes | Description

| 2621580
| Cannot specify options other than SVM name, comment and ipspace for a
Vserver that is being configured as the destination for SVM DR.

| 2621634
| "sync-source" SVM can only be created in a MetroCluster configuration.

| 2621657
| "sync-destination" SVM can only be created by the system.

```

| 13434884
| Cannot create an SVM because of incorrect fields.

| 13434885
| Non-UTF8 language(s) not supported.

| 13434888
| IPspace UUID and IPspace name mismatch.

| 13434889
| Internal Error. Wait and retry.

| 13434894
| Maximum allowed SVM jobs exceeded. Wait for the existing SVM jobs to complete and try again.

| 13434908
| Invalid SVM name. The name is already in use by another SVM, IPspace or cluster.

| 13434909
| Internal Error. Failed to identify the aggregate to host SVM root volume.

| 13434910
| Internal Error. Failed to allocate new SVM ID.

| 13434911
| Invalid SVM name. Maximum supported length is 41 if SVM is of type \"sync-source\", otherwise 47.

| 13434912
| Failed to find IPspace.

| 13434913
| Internal error: Failed to check if an SVM create operation is in progress. Contact technical support for assistance.

| 13434914
| Request to create the root volume of the SVM failed because there is not enough space in specified aggregate.

| 13434915
| Failed to unlock the SVM because SVM create or delete job is in progress. Wait a few minutes, and then try the command again.

| 13434916

```
| SVM is in the process of being created. Wait a few minutes, and then try  
the command again.
```

```
| 13434917  
| SVM creation successful.
```

```
| 13434918  
| IPspace name not provided for creating an SVM.  
|===
```

```
== Definitions
```

```
[.api-def-first-level]  
.See Definitions  
[%collapsible%closed]  
//Start collapsible Definitions block
```

```
=====
```

```
[#href]  
[.api-collapsible-fifth-title]  
href
```

```
[cols=3*,options=header]  
|===
```

```
|Name  
|Type  
|Description
```

```
|href  
|string  
a|
```

```
|===
```

```
[#_links]  
[.api-collapsible-fifth-title]  
_links  
[#aggregates]  
[.api-collapsible-fifth-title]  
aggregates
```

```
Aggregate
```

```
[cols=3*,options=header]  
|===
```



```
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#certificate]
[.api-collapsible-fifth-title]
certificate
```

Certificate for incoming TLS connection requests.

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID
```

```
|===
```

```
[#ad_domain]
[.api-collapsible-fifth-title]
ad_domain
```

```
[cols=3*,options=header]
|===
```

```

|Name
|Type
|Description

|fqdn
|string
a|The fully qualified domain name of the Windows Active Directory to which
this CIFS server belongs. A CIFS server appears as a member of Windows
server object in the Active Directory store.

|organizational_unit
|string
a|Specifies the organizational unit within the Active Directory domain to
associate with the CIFS server.

|password
|string
a|The account password used to add this CIFS server to the Active
Directory. This is not audited. Valid in POST only.

|user
|string
a|The user account used to add this CIFS server to the Active Directory.
Valid in POST only.

|===

[#cifs]
[.api-collapsible-fifth-title]
cifs

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_domain
|link:#ad_domain[ad_domain]
a|

|enabled

```

```
|boolean
a|Specifies whether or not the CIFS service is administratively enabled.
```

```
|name
|string
a|The NetBIOS name of the CIFS server.
```

```
|===
```

```
[#dns]
[.api-collapsible-fifth-title]
dns
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|domains
|array[string]
a|A list of DNS domains.
Domain names have the following requirements:
```

- * The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-", or "_".
- * The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.
- * The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.
- * The top level domain must contain only the following characters: A through Z, a through z.
- * The system reserves the following names:"all", "local", and "localhost".

```
|servers
|array[string]
a|The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
```

```
|===
```

```
[#node]
[.api-collapsible-fifth-title]
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#fc_port_reference]
[.api-collapsible-fifth-title]
fc_port_reference
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|The name of the FC port.
```

```
|node
|link:#node[node]
a|The node on which the FC port is located.
```

```
|uuid
```

```
|string
a|The unique identifier of the FC port.
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

The location of the Fibre Channel interface is defined by the location of its port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|port
|link:#fc_port_reference[fc_port_reference]
a|An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.
```

```
|===
```

```
[#fc_interface_svm]
[.api-collapsible-fifth-title]
fc_interface_svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|data_protocol
|string
a|The data protocol for which the Fibre Channel interface is configured.
```

```
|location
|link:#location[location]
```

a|The location of the Fibre Channel interface is defined by the location of its port.

|name

|string

a|The name of the Fibre Channel interface.

|uuid

|string

a|The unique identifier of the Fibre Channel interface.

|===

[#fcp]

[.api-collapsible-fifth-title]

fcp

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|Enable Fiber Channel Protocol (FCP)? Setting to true creates a service if not already created.

|===

[#ip]

[.api-collapsible-fifth-title]

ip

IP information

[cols=3*,options=header]

|===

|Name

|Type

```

|Description

|address
|string
a|IPv4 or IPv6 address


|netmask
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0).


|===

[#broadcast_domain]
[.api-collapsible-fifth-title]
broadcast_domain

Broadcast domain UUID along with a readable name.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the broadcast domain, scoped to its IPspace


|uuid
|string
a|Broadcast domain UUID


|===


[#home_node]
[.api-collapsible-fifth-title]
home_node

[cols=3*,options=header]
|===

```

```
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

Home_node is optional.

```
[cols=3*,options=header]
|===
```

```
|Name
|Type
|Description
```

```
|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name.
```

```
|home_node
|link:#home_node[home_node]
a|
```

```
|===
```

```
[#ip_interface_svm]
[.api-collapsible-fifth-title]
ip_interface_svm
```

Interface parameters. Name and home_node are optional.


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ip
|link:#ip[ip]
a|IP information

|location
|link:#location[location]
a|Home_node is optional.

|name
|string
a|The name of the interface (optional).

|service_policy
|string
a|Built-in service policies for SVMs.

|services
|array[string]
a|The services associated with the interface.

|uuid
|string
a|The UUID that uniquely identifies the interface.

|===

[#ipspace]
[.api-collapsible-fifth-title]
ipspace

Either the UUID or name may be supplied on input.

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|name
|string
a|IPspace name


|uuid
|string
a|IPspace UUID


|===


[#iscsi]
[.api-collapsible-fifth-title]
iscsi

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Enable iSCSI? Setting to true creates a service if not already created.


|===


[#ldap]
[.api-collapsible-fifth-title]
ldap

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_domain

```

```

|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST.

|base_dn
|string
a|Specifies the default base DN for all searches.

|bind_dn
|string
a|Specifies the user that binds to the LDAP servers. SVM API supports
anonymous binding. For Simple and SASL LDAP binding, use the LDAP API
endpoint.

|enabled
|boolean
a|Enable LDAP? Setting to true creates a configuration if not already
created.

|servers
|array[string]
a|

|===

[#nfs]
[.api-collapsible-fifth-title]
nfs

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Enable NFS? Setting to true creates a service if not already created.

|===

```

```

[#nis]
[.api-collapsible-fifth-title]
nis

[cols=3*,options=header]
|===
|Name
|Type
|Description

|domain
|string
a|The NIS domain to which this configuration belongs.

|enabled
|boolean
a|Enable NIS? Setting to true creates a configuration if not already
created.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|===

[#nsswitch]
[.api-collapsible-fifth-title]
nsswitch

Name service switch configuration

[cols=3*,options=header]
|===
|Name
|Type
|Description

|group
|array[string]

```

```
a|Group sources
```

```
|hosts
```

```
|array[string]
```

```
a|Host sources
```

```
|namemap
```

```
|array[string]
```

```
a|NameMap sources
```

```
|netgroup
```

```
|array[string]
```

```
a|NetGroup sources
```

```
|passwd
```

```
|array[string]
```

```
a|Password sources
```

```
|===
```

```
[#nvme]
```

```
[.api-collapsible-fifth-title]
```

```
nvme
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|Enable NVMe? Setting to true creates a service if not already created.
```

```
|===
```

```
[#ip_info]
```

```
[.api-collapsible-fifth-title]
```

```
ip_info
```

IP information

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|address
```

```
|string
```

```
a|IPv4 or IPv6 address
```

```
|netmask
```

```
|string
```

```
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you  
must set the netmask length. The default value is 64. Output is always  
netmask length.
```

```
|===
```

```
[#network_route_for_svm]
```

```
[.api-collapsible-fifth-title]
```

```
network_route_for_svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|destination
```

```
|link:#ip_info[ip_info]
```

```
a|IP information
```

```
|gateway
```

```
|string
```

```
a|The IP address of the gateway router leading to the destination.
```

```
|===
```

```
[#certificate]
[.api-collapsible-fifth-title]
certificate
```

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID
```

```
|===
```

```
[#s3]
[.api-collapsible-fifth-title]
s3
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|certificate
|link:#certificate[certificate]
a|Specifies the certificate that will be used for creating HTTPS
connections to the S3 server.
```

```
|enabled
|boolean
a|Specifies whether or not to enable S3. Setting this value to true
```

creates a service if one is not yet created.

|is_http_enabled

|boolean

a|Specifies whether HTTP is enabled on the S3 server. By default, HTTP is disabled on the S3 server.

|is_https_enabled

|boolean

a|Specifies whether HTTPS is enabled on the S3 server. By default, HTTPS is enabled on the S3 server.

|name

|string

a|Specifies the name of the S3 server. A server name length can range from 1 to 15 characters and can only contain the following combination of characters 0-9, A-Z, a-z, ".", and "-".

|port

|integer

a|Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80.

|secure_port

|integer

a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443.

|===

[#snapmirror]

[.api-collapsible-fifth-title]

snapmirror

Specifies attributes for SVM DR protection.

[cols=3*,options=header]

|===

|Name


```

|Type
|Description

|is_protected
|boolean
a|Specifies whether the SVM is a SnapMirror source SVM, using SnapMirror
to protect its data.

|protected_volumes_count
|integer
a|Specifies the number of SVM DR protected volumes in the SVM.

|===

[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

This is a reference to the Snapshot copy policy.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#volume_efficiency_policy]
[.api-collapsible-fifth-title]
volume_efficiency_policy

This is a reference to the volume efficiency policy.

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the volume efficiency policy

|uuid
|string
a|Unique identifier of the volume efficiency policy.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|array[link:#aggregates[aggregates]]
a|List of allowed aggregates for SVM volumes. An administrator is allowed
to create volumes on these aggregates.

|aggregates_delegated
|boolean
a|This property is true when the administrator has delegated the
aggregates for the SVM volumes.

|cifs
|link:#cifs[cifs]
a|

|comment

```

```

|string
a|Comment

|dns
|link:#dns[dns]
a|

|fc_interfaces
|array[link:#fc_interface_svm[fc_interface_svm]]
a|FC Interface for the SVM

|fc
|link:#fc[fc]
a|

|ip_interfaces
|array[link:#ip_interface_svm[ip_interface_svm]]
a|IP interfaces for the SVM

|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name may be supplied on input.

|iscsi
|link:#iscsi[iscsi]
a|

|language
|string
a|Default volume language code. UTF-8 encoded languages are valid in POST
or PATCH. Non UTF-8 language encodings are for backward compatibility and
are not valid input for POST and PATCH requests.

|ldap
|link:#ldap[ldap]
a|

|name
|string
a|The name of the SVM.

```

```

|nfs
|link:#nfs[nfs]
a|

|nis
|link:#nis[nis]
a|

|nsswitch
|link:#nsswitch[nsswitch]
a|Name service switch configuration

|nvme
|link:#nvme[nvme]
a|

|routes
|array[link:#network_route_for_svm[network_route_for_svm]]
a|Optional array of routes for the SVM

|s3
|link:#s3[s3]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SVM DR protection.

|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the Snapshot copy policy.

|subtype
|string
a|SVM subtype. The SVM subtype sync_destination is created automatically
when an SVM of subtype sync_source is created on the source MetroCluster
cluster. A POST request with sync_destination as SVM subtype is invalid.

|uuid
|string
a|The unique identifier of the SVM.

```

```
|volume_efficiency_policy
|link:#volume_efficiency_policy[volume_efficiency_policy]
a|This is a reference to the volume efficiency policy.
```

```
|===
```

```
[#job_link]
[.api-collapsible-fifth-title]
job_link
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|uuid
|string
```

```
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|code
|string
```

```
a|Argument code
```

```
|message
|string
```

```
a|Message argument
```

```

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[ID60aa8599950766f765dc53c87ef47197]]
= Delete an SVM

[.api-doc-operation .api-doc-operation-delete]#DELETE# [.api-doc-code-
block]#`/svm/svms/{uuid}`#

```

Introduced In: 9.6

Deletes an SVM. As a prerequisite, SVM objects must be deleted first. SnapMirror relationships must be deleted and data volumes must be offline and deleted.

* The number of parallel SVMs that can be created must not be greater than five.

* If a sixth SVM POST request is issued, the following error message is generated: "Maximum allowed SVM jobs exceeded. Wait for the existing SVM jobs to complete and try again."

== Related ONTAP commands

* `vserver delete`

== Example

Deleting an individual SVM in the cluster.

```
-----  
DELETE "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"  
-----
```

== Learn more

* xref:{relative_path}svm_svms_endpoint_overview.html[DOC /svm/svms]

== Parameters

[cols=5*,options=header]

|===

|Name

|Type

|In

|Required

|Description

|uuid

|string

|path

|True

a|Filter by UUID

```
|return_timeout  
|integer  
|query  
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

```
* Default value: 1  
* Max value: 120  
* Min value: 0
```

```
|===
```

```
== Response
```

Status: 202, Accepted


```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
====

== Error

```

Status: Default

ONTAP Error Response Codes

```

|===
| Error codes | Description

| 13434894
| Maximum allowed SVM jobs exceeded. Wait and retry.
|===

```

```

[cols=3*,options=header]

```

```

|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
=====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
=====

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
=====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type

```

```

|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#job_link]
[.api-collapsible-fifth-title]
job_link

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|uuid
|string
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.

|===

```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```
|message
```

```
|string
```

```
a|Message argument
```

```
|===
```

```
[#error]
```

```
[.api-collapsible-fifth-title]
```

```
error
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|arguments
```

```
|array[link:#error_arguments[error_arguments]]
```

```
a|Message arguments
```

```
|code
```

```
|string
```

```
a|Error code
```

```
|message
```

```
|string
```

a|Error message

|target

|string

a|The target parameter that caused the error.

|===

//end collapsible .Definitions block

====

[[ID00f63d3c67d240d8e0093246dee09b44]]

= Retrieve SVM properties

[.api-doc-operation .api-doc-operation-get]#GET# [.api-doc-code-block]#`/svm/svms/{uuid}`#

Introduced In: 9.6

Retrieves the properties for an individual SVM. This includes protocol configurations such as CIFS and NFS, export policies, name service configurations, and network services.

== Important notes

* The SVM object includes a large set of fields and can be expensive to retrieve.

* REST APIs only expose a data SVM as an SVM.

== Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [xref:{relative_path}getting_started_with_the_ontap_rest_api.html#Requesting_specific_fields\[Requesting specific fields\]](#) to learn more.

* `snapmirror.+*+`

== Example

Retrieving an individual SVM in the cluster

```
----  
GET "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"  
----
```

== Parameters

```
[cols=5*,options=header]  
|===  
  
|Name  
|Type  
|In  
|Required  
|Description  
  
|uuid  
|string  
|path  
|True  
a|Filter by UUID  
  
|fields  
|array[string]  
|query  
|False  
a|Specify the fields to return.  
  
|===
```

== Response

Status: 200, Ok

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description  
  
|_links  
|link:#_links[_links]  
a|
```

```
|aggregates
|array[link:#aggregates[aggregates]]
a|List of allowed aggregates for SVM volumes. An administrator is allowed
to create volumes on these aggregates.
```

```
|aggregates_delegated
|boolean
a|This property is true when the administrator has delegated the
aggregates for the SVM volumes.
```

```
|certificate
|link:#certificate[certificate]
a|Certificate for incoming TLS connection requests.
```

```
|cifs
|link:#cifs[cifs]
a|
```

```
|comment
|string
a|Comment
```

```
|dns
|link:#dns[dns]
a|
```

```
|fc_interfaces
|array[link:#fc_interface_svm[fc_interface_svm]]
a|FC Interface for the SVM
```

```
|fc
|link:#fc[fc]
a|
```

```
|ip_interfaces
|array[link:#ip_interface_svm[ip_interface_svm]]
a|IP interfaces for the SVM
```

```
|ipspace
|link:#ipspace[ipspace]
a|Either the UUID or name may be supplied on input.
```

```

|iscsi
|link:#iscsi[iscsi]
a|

|language
|string
a|Default volume language code. UTF-8 encoded languages are valid in POST
or PATCH. Non UTF-8 language encodings are for backward compatibility and
are not valid input for POST and PATCH requests.

|ldap
|link:#ldap[ldap]
a|

|name
|string
a|The name of the SVM.

|nfs
|link:#nfs[nfs]
a|

|nis
|link:#nis[nis]
a|

|nsswitch
|link:#nsswitch[nsswitch]
a|Name service switch configuration

|nvme
|link:#nvme[nvme]
a|

|s3
|link:#s3[s3]
a|

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SVM DR protection.

```



```
|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the Snapshot copy policy.
```

```
|state
|string
a|SVM State
```

```
|subtype
|string
a|SVM subtype. The SVM subtype sync_destination is created automatically
when an SVM of subtype sync_source is created on the source MetroCluster
cluster. A POST request with sync_destination as SVM subtype is invalid.
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|volume_efficiency_policy
|link:#volume_efficiency_policy[volume_efficiency_policy]
a|This is a reference to the volume efficiency policy.
```

```
|===
```

```
.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  ]
}
```

```

    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
],
"certificate": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "cert1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"cifs": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ad_domain": {
    "fqdn": "example.com",
    "organizational_unit": "string"
  },
  "name": "CIFS1"
},
"comment": "string",
"dns": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "domains": [
    "example.com",
    "example2.example3.com"
  ],
  "servers": [
    "10.224.65.20",
    "2001:db08:a0b:12f0::1"
  ]
},
"fc_interfaces": [
  {
    "_links": {
      "self": {

```

```

        "href": "/api/resourcelink"
    }
},
"data_protocol": "string",
"name": "lif1",
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
],
"fcg": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    }
},
"ip_interfaces": [
    {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "ip": {
            "address": "10.10.10.7"
        },
        "name": "lif1",
        "services": [
            "data_nfs"
        ],
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
],
"ipspace": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "exchange",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"iscsi": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    }
}

```

```

    }
  },
  "language": "c.utf_8",
  "ldap": {
    "ad_domain": "string",
    "base_dn": "string",
    "bind_dn": "string",
    "servers": [
      "string"
    ]
  },
  "name": "svml",
  "nfs": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "nis": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "domain": "string",
  "servers": [
    "string"
  ]
},
"nsswitch": {
  "group": [
    "string"
  ],
  "hosts": [
    "string"
  ],
  "namemap": [
    "string"
  ],
  "netgroup": [
    "string"
  ],
  "passwd": [
    "string"
  ]
}

```

```

},
"nvme": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"s3": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "s3-server-1"
},
"snapmirror": {
  "protected_volumes_count": 0
},
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"state": "running",
"subtype": "string",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7",
"volume_efficiency_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},

```

```

    "name": "default",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Error

```

Status: Default, Error

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|error
|link:#error[error]
a|

|===

.Example error
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
====

== Definitions

[.api-def-first-level]
.See Definitions

```

```

[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links

[cols=3*,options=header]
|===
|Name
|Type
|Description

|self
|link:#href[href]
a|

|===

[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name

```

```
|Type
|Description

|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#certificate]
[.api-collapsible-fifth-title]
certificate
```

Certificate for incoming TLS connection requests.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID
```

```
|===
```



```

[#ad_domain]
[.api-collapsible-fifth-title]
ad_domain

[cols=3*,options=header]
|===
|Name
|Type
|Description

|fqdn
|string
a|The fully qualified domain name of the Windows Active Directory to which
this CIFS server belongs. A CIFS server appears as a member of Windows
server object in the Active Directory store.

|organizational_unit
|string
a|Specifies the organizational unit within the Active Directory domain to
associate with the CIFS server.

|===

[#cifs]
[.api-collapsible-fifth-title]
cifs

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ad_domain
|link:#ad_domain[ad_domain]
a|

|enabled
|boolean

```

a|Specifies whether or not the CIFS service is administratively enabled.

|name

|string

a|The NetBIOS name of the CIFS server.

|===

[#dns]

[.api-collapsible-fifth-title]

dns

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|domains

|array[string]

a|A list of DNS domains.

Domain names have the following requirements:

- * The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-" or "_".

- * The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.

- * The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.

- * The top level domain must contain only the following characters: A through Z, a through z.

- * The system reserves the following names:"all", "local", and "localhost".

|servers

|array[string]

a|The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.

|===

[#node]

[.api-collapsible-fifth-title]

node

The node on which the FC port is located.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|The name of the node on which the FC port is located.

|===

[#fc_port_reference]

[.api-collapsible-fifth-title]

fc_port_reference

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|The name of the FC port.

```
|node
|link:#node[node]
a|The node on which the FC port is located.
```

```
|uuid
|string
a|The unique identifier of the FC port.
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

The location of the Fibre Channel interface is defined by the location of its port.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|port
|link:#fc_port_reference[fc_port_reference]
a|An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.
```

```
|===
```

```
[#fc_interface_svm]
[.api-collapsible-fifth-title]
fc_interface_svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```

|_links
|link:#_links[_links]
a|

|data_protocol
|string
a|The data protocol for which the Fibre Channel interface is configured.

|name
|string
a|The name of the Fibre Channel interface.

|uuid
|string
a|The unique identifier of the Fibre Channel interface.

|===

[#fcp]
[.api-collapsible-fifth-title]
fcp

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|Enable Fiber Channel Protocol (FCP)? Setting to true creates a service
if not already created.

|===

[#ip]
[.api-collapsible-fifth-title]

```

ip

IP information

[cols=3*,options=header]

|===

|Name

|Type

|Description

|address

|string

a|IPv4 or IPv6 address

|===

[#broadcast_domain]

[.api-collapsible-fifth-title]

broadcast_domain

Broadcast domain UUID along with a readable name.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|_links

|link:#_links[_links]

a|

|name

|string

a|Name of the broadcast domain, scoped to its IPspace

|uuid

|string

a|Broadcast domain UUID

|===

```
[#home_node]
[.api-collapsible-fifth-title]
home_node
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

Home_node is optional.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name.
```

```
|home_node
|link:#home_node[home_node]
```

```

a|

|===

[#ip_interface_svm]
[.api-collapsible-fifth-title]
ip_interface_svm

Interface parameters.  Name and home_node are optional.


[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|ip
|link:#ip[ip]
a|IP information


|name
|string
a|The name of the interface (optional).


|services
|array[string]
a|The services associated with the interface.


|uuid
|string
a|The UUID that uniquely identifies the interface.


|===

[#ipspace]
[.api-collapsible-fifth-title]

```


ipspace

Either the UUID or name may be supplied on input.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|IPspace name
```

```
|uuid
```

```
|string
```

```
a|IPspace UUID
```

```
|===
```

```
[#iscsi]
```

```
[.api-collapsible-fifth-title]
```

```
iscsi
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|enabled
```

```
|boolean
```

```
a|Enable iSCSI? Setting to true creates a service if not already created.
```

```

|===

[#ldap]
[.api-collapsible-fifth-title]
ldap

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_domain
|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST.

|base_dn
|string
a|Specifies the default base DN for all searches.

|bind_dn
|string
a|Specifies the user that binds to the LDAP servers. SVM API supports
anonymous binding. For Simple and SASL LDAP binding, use the LDAP API
endpoint.

|enabled
|boolean
a|Enable LDAP? Setting to true creates a configuration if not already
created.

|servers
|array[string]
a|

|===

[#nfs]
[.api-collapsible-fifth-title]

```

```

nfs

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|Enable NFS? Setting to true creates a service if not already created.

|===

[#nis]
[.api-collapsible-fifth-title]
nis

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|domain
|string
a|The NIS domain to which this configuration belongs.

|enabled
|boolean
a|Enable NIS? Setting to true creates a configuration if not already
created.

|servers
|array[string]

```

a|A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

|===

```
[#nsswitch]
[.api-collapsible-fifth-title]
nsswitch
```

Name service switch configuration

```
[cols=3*,options=header]
```

|===

```
|Name
|Type
|Description
```

```
|group
|array[string]
a|Group sources
```

```
|hosts
|array[string]
a|Host sources
```

```
|namemap
|array[string]
a|NameMap sources
```

```
|netgroup
|array[string]
a|NetGroup sources
```

```
|passwd
|array[string]
a|Password sources
```

|===

```

[#nvme]
[.api-collapsible-fifth-title]
nvme

[cols=3*,options=header]
|===
|Name
|Type
|Description

|_links
|link:#_links[_links]
a|

|enabled
|boolean
a|Enable NVMe? Setting to true creates a service if not already created.

|===

[#ip_info]
[.api-collapsible-fifth-title]
ip_info

IP information

[cols=3*,options=header]
|===
|Name
|Type
|Description

|address
|string
a|IPv4 or IPv6 address

|family
|string
a|IPv4 or IPv6

|netmask

```

```
|string
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you
must set the netmask length. The default value is 64. Output is always
netmask length.
```

```
|===
```

```
[#network_route_for_svm]
[.api-collapsible-fifth-title]
network_route_for_svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|destination
```

```
|link:#ip_info[ip_info]
```

```
a|IP information
```

```
|gateway
```

```
|string
```

```
a|The IP address of the gateway router leading to the destination.
```

```
|===
```

```
[#certificate]
```

```
[.api-collapsible-fifth-title]
```

```
certificate
```

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|name
```

```
|string
```

```
a|Certificate name
```

```
|uuid
```

```
|string
```

```
a|Certificate UUID
```

```
|===
```

```
[#s3]
```

```
[.api-collapsible-fifth-title]
```

```
s3
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|_links
```

```
|link:#_links[_links]
```

```
a|
```

```
|certificate
```

```
|link:#certificate[certificate]
```

```
a|Specifies the certificate that will be used for creating HTTPS  
connections to the S3 server.
```

```
|enabled
```

```
|boolean
```

```
a|Specifies whether or not to enable S3. Setting this value to true  
creates a service if one is not yet created.
```

```
|is_http_enabled
```

```
|boolean
```

```
a|Specifies whether HTTP is enabled on the S3 server. By default, HTTP is  
disabled on the S3 server.
```

|is_https_enabled

|boolean

a|Specifies whether HTTPS is enabled on the S3 server. By default, HTTPS is enabled on the S3 server.

|name

|string

a|Specifies the name of the S3 server. A server name length can range from 1 to 15 characters and can only contain the following combination of characters 0-9, A-Z, a-z, ".", and "-".

|port

|integer

a|Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80.

|secure_port

|integer

a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443.

|===

[#snapmirror]

[.api-collapsible-fifth-title]

snapmirror

Specifies attributes for SVM DR protection.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|is_protected

|boolean

a|Specifies whether the SVM is a SnapMirror source SVM, using SnapMirror to protect its data.


```
|protected_volumes_count
|integer
a|Specifies the number of SVM DR protected volumes in the SVM.
```

```
|===
```

```
[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy
```

This is a reference to the Snapshot copy policy.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|_links
|link:#_links[_links]
a|
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#volume_efficiency_policy]
[.api-collapsible-fifth-title]
volume_efficiency_policy
```

This is a reference to the volume efficiency policy.

```
[cols=3*,options=header]
|===
|Name
```

```

|Type
|Description

|_links
|link:#_links[_links]
a|

|name
|string
a|Name of the volume efficiency policy

|uuid
|string
a|Unique identifier of the volume efficiency policy.

|===

[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments

[cols=3*,options=header]
|===
|Name
|Type
|Description

|code
|string
a|Argument code

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

[[IDb3218f82ac803a2ac720b8d3c6a0c165]]
= Update SVM properties

[.api-doc-operation .api-doc-operation-patch]#PATCH# [.api-doc-code-
block]#`/svm/svms/{uuid}`#

*Introduced In:* 9.6

Updates one or more of the following properties of an individual SVM: SVM
name, SVM default volume language code, SVM comment, and SVM state.

== Related ONTAP commands

```

```
* `vserver modify`  
* `vserver rename`  
* `vserver start`  
* `vserver stop`  
* `security ssl modify`
```

== Examples

. Stops an SVM and updates the "comment" field for an individual SVM

```
-----  
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"  
'{"state":"stopped", "comment":"This SVM is stopped."}'  
-----
```

. Starts an SVM and updates the "comment" field for an individual SVM

```
-----  
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"  
'{"state":"running", "comment":"This SVM is running."}'  
-----
```

. Updates the "language" field for an individual SVM

```
-----  
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"  
'{"language":"en.UTF-8"}'  
-----
```

. Updates the "name" field for an SVM or renames the SVM

```
-----  
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"  
'{"name":"svm_new"}'  
-----
```

. Updates the aggregates for an individual SVM

```
-----  
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"  
'{"aggregates":{"name":["aggr1","aggr2","aggr3"]}}'  
-----
```

. Updates the Snapshot copy policy for an individual SVM

```
-----  
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
```

```
'{"snapshot_policy":{"name":"custom1"}}'
```

```
----
```

. Updates the TLS certificate for an individual SVM

```
----
```

```
PATCH "/api/svm/svms/f16f0935-5281-11e8-b94d-005056b46485"
'{"certificate":{"uuid":"1cd8a442-86d1-11e0-ae1c-123478563412"}}'
```

```
----
```

== Learn more

* xref:{relative_path}svm_svms_endpoint_overview.html[DOC /svm/svms]

== Parameters

```
[cols=5*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|In
```

```
|Required
```

```
|Description
```

```
|uuid
```

```
|string
```

```
|path
```

```
|True
```

```
a|Filter by UUID
```

```
|return_timeout
```

```
|integer
```

```
|query
```

```
|False
```

a|The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.

* Default value: 1

* Max value: 120

* Min value: 0

|===

== Request Body

[cols=3*,options=header]

|===

|Name

|Type

|Description

|aggregates

|array[link:#aggregates[aggregates]]

a|List of allowed aggregates for SVM volumes. An administrator is allowed to create volumes on these aggregates.

|aggregates_delegated

|boolean

a|This property is true when the administrator has delegated the aggregates for the SVM volumes.

|certificate

|link:#certificate[certificate]

a|Certificate for incoming TLS connection requests.

|comment

|string

a|Comment

|fc_interfaces

|array[link:#fc_interface_svm[fc_interface_svm]]

a|FC Interface for the SVM

|ip_interfaces

|array[link:#ip_interface_svm[ip_interface_svm]]

a|IP interfaces for the SVM

|language

```
|string
a|Default volume language code. UTF-8 encoded languages are valid in POST
or PATCH. Non UTF-8 language encodings are for backward compatibility and
are not valid input for POST and PATCH requests.
```

```
|name
|string
a|The name of the SVM.
```

```
|nsswitch
|link:#nsswitch[nsswitch]
a|Name service switch configuration
```

```
|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SVM DR protection.
```

```
|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the Snapshot copy policy.
```

```
|state
|string
a|SVM State
```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|volume_efficiency_policy
|link:#volume_efficiency_policy[volume_efficiency_policy]
a|This is a reference to the volume efficiency policy.
```

```
|===
```

```
.Example request
[%collapsible%closed]
=====
```

```
[source,json,subs=+macros]
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "certificate": {
    "name": "cert1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "comment": "string",
  "fc_interfaces": [
    {
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "ip_interfaces": [
    {
      "services": [
        "data_nfs"
      ],
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "language": "c.utf_8",
  "name": "svm1",
  "nsswitch": {
    "group": [
      "string"
    ],
    "hosts": [
      "string"
    ],
    "namemap": [
      "string"
    ],
    "netgroup": [
      "string"
    ],
    "passwd": [
      "string"
    ]
  },
  "snapmirror": {
```



```

    "protected_volumes_count": 0
  },
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "running",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7",
  "volume_efficiency_policy": {
    "name": "default",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
====

== Response

```

Status: 202, Accepted

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|job
|link:#job_link[job_link]
a|

|===

.Example response
[%collapsible%closed]
====
[source,json,subs=+macros]
{
  "job": {
    "uuid": "string"
  }
}
====

== Error

```

ONTAP Error Response Codes

|===

| Error codes | Description

| 13434880

| Failed to modify SVM parameters.

| 13434881

| Failed to rename SVM.

| 13434883

| SVM parameters except name modified successfully.

| 13434885

| Non-UTF8 language(s) not supported.

| 13434886

| Invalid Snapshot copy policy.

| 13434902

| Modification of NSSwitch parameters failed for the SVM.

| 13434906

| Operation not supported for an SVM of type sync-destination.

| 12451843

| Certificate does not exist.

| 13434908

| Invalid SVM name. The name is already in use by another SVM, IPspace or cluster.

| 13434916

| SVM is in the process of being created. Wait a few minutes, and then try the command again.

| 13434915

| Failed to unlock the SVM because SVM create or delete job is in progress. Wait a few minutes, and then try the command again.

| 13434911

| Invalid SVM name. Maximum supported length is 41 if SVM is of type \"sync-source\", otherwise 47.

```

|===

== Definitions

[.api-def-first-level]
.See Definitions
[%collapsible%closed]
//Start collapsible Definitions block
====
[#href]
[.api-collapsible-fifth-title]
href

[cols=3*,options=header]
|===
|Name
|Type
|Description

|href
|string
a|

|===

[#_links]
[.api-collapsible-fifth-title]
_links
[#aggregates]
[.api-collapsible-fifth-title]
aggregates

Aggregate

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

```

```
|uuid
|string
a|
```

```
|===
```

```
[#certificate]
[.api-collapsible-fifth-title]
certificate
```

Certificate for incoming TLS connection requests.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID
```

```
|===
```

```
[#ad_domain]
[.api-collapsible-fifth-title]
ad_domain
[#cifs]
[.api-collapsible-fifth-title]
cifs
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|ad_domain
```

```
|link:#ad_domain[ad_domain]
```

```
a|
```

```
|enabled
```

```
|boolean
```

```
a|Specifies whether or not the CIFS service is administratively enabled.
```

```
|name
```

```
|string
```

```
a|The NetBIOS name of the CIFS server.
```

```
|===
```

```
[#dns]
```

```
[.api-collapsible-fifth-title]
```

```
dns
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|domains
```

```
|array[string]
```

```
a|A list of DNS domains.
```

```
Domain names have the following requirements:
```

- * The name must contain only the following characters: A through Z, a through z, 0 through 9, ".", "-" or "_".

- * The first character of each label, delimited by ".", must be one of the following characters: A through Z or a through z or 0 through 9.

- * The last character of each label, delimited by ".", must be one of the following characters: A through Z, a through z, or 0 through 9.

- * The top level domain must contain only the following characters: A through Z, a through z.

- * The system reserves the following names:"all", "local", and "localhost".

```
|servers
```

```
|array[string]
```

```
a|The list of IP addresses of the DNS servers. Addresses can be either IPv4 or IPv6 addresses.
```

```
|===
```

```
[#node]  
[.api-collapsible-fifth-title]  
node
```

The node on which the FC port is located.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|name  
|string  
a|The name of the node on which the FC port is located.
```

```
|===
```

```
[#fc_port_reference]  
[.api-collapsible-fifth-title]  
fc_port_reference
```

An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

```
[cols=3*,options=header]  
|===  
|Name  
|Type  
|Description
```

```
|name  
|string  
a|The name of the FC port.
```

```
|node  
|link:#node[node]
```

a|The node on which the FC port is located.

|uuid

|string

a|The unique identifier of the FC port.

|===

[#location]

[.api-collapsible-fifth-title]

location

The location of the Fibre Channel interface is defined by the location of its port.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|port

|link:#fc_port_reference[fc_port_reference]

a|An FC port is the physical port of an FC adapter on a cluster node that can be connected to an FC network.

|===

[#fc_interface_svm]

[.api-collapsible-fifth-title]

fc_interface_svm

[cols=3*,options=header]

|===

|Name

|Type

|Description

|uuid

|string

a|The unique identifier of the Fibre Channel interface.

|===

[#fcp]

[.api-collapsible-fifth-title]

fcp

[cols=3*,options=header]

|===

|Name

|Type

|Description

|enabled

|boolean

a|Enable Fiber Channel Protocol (FCP)? Setting to true creates a service if not already created.

|===

[#ip]

[.api-collapsible-fifth-title]

ip

IP information

[#broadcast_domain]

[.api-collapsible-fifth-title]

broadcast_domain

Broadcast domain UUID along with a readable name.

[cols=3*,options=header]

|===

|Name

|Type

|Description

|name

|string

a|Name of the broadcast domain, scoped to its IPspace


```
|uuid
|string
a|Broadcast domain UUID
```

```
|===
```

```
[#home_node]
[.api-collapsible-fifth-title]
home_node
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|name
|string
a|
```

```
|uuid
|string
a|
```

```
|===
```

```
[#location]
[.api-collapsible-fifth-title]
location
```

Home_node is optional.

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
|Type
|Description
```

```
|broadcast_domain
|link:#broadcast_domain[broadcast_domain]
a|Broadcast domain UUID along with a readable name.
```

```
|home_node
|link:#home_node[home_node]
a|
```

```
|===
```

```
[#ip_interface_svm]
[.api-collapsible-fifth-title]
ip_interface_svm
```

Interface parameters. Name and home_node are optional.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|ip
|link:#ip[ip]
a|IP information
```

```
|services
|array[string]
a|The services associated with the interface.
```

```
|uuid
|string
a|The UUID that uniquely identifies the interface.
```

```
|===
```

```
[#ipspace]
[.api-collapsible-fifth-title]
ipspace
```

Either the UUID or name may be supplied on input.

```
[cols=3*,options=header]
```

```

|===
|Name
|Type
|Description

|name
|string
a|IPspace name


|uuid
|string
a|IPspace UUID


|===


[#iscsi]
[.api-collapsible-fifth-title]
iscsi

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Enable iSCSI? Setting to true creates a service if not already created.


|===


[#ldap]
[.api-collapsible-fifth-title]
ldap

[cols=3*,options=header]
|===
|Name
|Type
|Description

|ad_domain

```

```

|string
a|This parameter specifies the name of the Active Directory domain
used to discover LDAP servers for use by this client.
This is mutually exclusive with `servers` during POST.

|base_dn
|string
a|Specifies the default base DN for all searches.

|bind_dn
|string
a|Specifies the user that binds to the LDAP servers. SVM API supports
anonymous binding. For Simple and SASL LDAP binding, use the LDAP API
endpoint.

|enabled
|boolean
a|Enable LDAP? Setting to true creates a configuration if not already
created.

|servers
|array[string]
a|

|===

[#nfs]
[.api-collapsible-fifth-title]
nfs

[cols=3*,options=header]
|===
|Name
|Type
|Description

|enabled
|boolean
a|Enable NFS? Setting to true creates a service if not already created.

|===

```

```

[#nis]
[.api-collapsible-fifth-title]
nis

[cols=3*,options=header]
|===
|Name
|Type
|Description

|domain
|string
a|The NIS domain to which this configuration belongs.

|enabled
|boolean
a|Enable NIS? Setting to true creates a configuration if not already
created.

|servers
|array[string]
a|A list of hostnames or IP addresses of NIS servers used
by the NIS domain configuration.

|===

[#nsswitch]
[.api-collapsible-fifth-title]
nsswitch

Name service switch configuration

[cols=3*,options=header]
|===
|Name
|Type
|Description

|group
|array[string]

```

```
a|Group sources
```

```
|hosts
```

```
|array[string]
```

```
a|Host sources
```

```
|namemap
```

```
|array[string]
```

```
a|NameMap sources
```

```
|netgroup
```

```
|array[string]
```

```
a|NetGroup sources
```

```
|passwd
```

```
|array[string]
```

```
a|Password sources
```

```
|===
```

```
[#nvme]
```

```
[.api-collapsible-fifth-title]
```

```
nvme
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|enabled
```

```
|boolean
```

```
a|Enable NVMe? Setting to true creates a service if not already created.
```

```
|===
```

```
[#ip_info]
```

```
[.api-collapsible-fifth-title]
```

```
ip_info
```

IP information

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|address
```

```
|string
```

```
a|IPv4 or IPv6 address
```

```
|netmask
```

```
|string
```

```
a|Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you  
must set the netmask length. The default value is 64. Output is always  
netmask length.
```

```
|===
```

```
[#network_route_for_svm]
```

```
[.api-collapsible-fifth-title]
```

```
network_route_for_svm
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|destination
```

```
|link:#ip_info[ip_info]
```

```
a|IP information
```

```
|gateway
```

```
|string
```

```
a|The IP address of the gateway router leading to the destination.
```

```
|===
```

```
[#certificate]
[.api-collapsible-fifth-title]
certificate
```

Specifies the certificate that will be used for creating HTTPS connections to the S3 server.

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|name
|string
a|Certificate name
```

```
|uuid
|string
a|Certificate UUID
```

```
|===
```

```
[#s3]
[.api-collapsible-fifth-title]
s3
```

```
[cols=3*,options=header]
|===
|Name
|Type
|Description
```

```
|certificate
|link:#certificate[certificate]
a|Specifies the certificate that will be used for creating HTTPS
connections to the S3 server.
```

```
|enabled
|boolean
a|Specifies whether or not to enable S3. Setting this value to true
```


creates a service if one is not yet created.

|is_http_enabled

|boolean

a|Specifies whether HTTP is enabled on the S3 server. By default, HTTP is disabled on the S3 server.

|is_https_enabled

|boolean

a|Specifies whether HTTPS is enabled on the S3 server. By default, HTTPS is enabled on the S3 server.

|name

|string

a|Specifies the name of the S3 server. A server name length can range from 1 to 15 characters and can only contain the following combination of characters 0-9, A-Z, a-z, ".", and "-".

|port

|integer

a|Specifies the HTTP listener port for the S3 server. By default, HTTP is enabled on port 80.

|secure_port

|integer

a|Specifies the HTTPS listener port for the S3 server. By default, HTTPS is enabled on port 443.

|===

[#snapmirror]

[.api-collapsible-fifth-title]

snapmirror

Specifies attributes for SVM DR protection.

[cols=3*,options=header]

|===

|Name

```

|Type
|Description

|is_protected
|boolean
a|Specifies whether the SVM is a SnapMirror source SVM, using SnapMirror
to protect its data.

|protected_volumes_count
|integer
a|Specifies the number of SVM DR protected volumes in the SVM.

|===

[#snapshot_policy]
[.api-collapsible-fifth-title]
snapshot_policy

This is a reference to the Snapshot copy policy.

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|

|uuid
|string
a|

|===

[#volume_efficiency_policy]
[.api-collapsible-fifth-title]
volume_efficiency_policy

This is a reference to the volume efficiency policy.

```

```

[cols=3*,options=header]
|===
|Name
|Type
|Description

|name
|string
a|Name of the volume efficiency policy

|uuid
|string
a|Unique identifier of the volume efficiency policy.

|===

[#svm]
[.api-collapsible-fifth-title]
svm

[cols=3*,options=header]
|===
|Name
|Type
|Description

|aggregates
|array[link:#aggregates[aggregates]]
a|List of allowed aggregates for SVM volumes. An administrator is allowed
to create volumes on these aggregates.

|aggregates_delegated
|boolean
a|This property is true when the administrator has delegated the
aggregates for the SVM volumes.

|certificate
|link:#certificate[certificate]
a|Certificate for incoming TLS connection requests.

```

```

|comment
|string
a|Comment

|fc_interfaces
|array[link:#fc_interface_svm[fc_interface_svm]]
a|FC Interface for the SVM

|ip_interfaces
|array[link:#ip_interface_svm[ip_interface_svm]]
a|IP interfaces for the SVM

|language
|string
a|Default volume language code. UTF-8 encoded languages are valid in POST
or PATCH. Non UTF-8 language encodings are for backward compatibility and
are not valid input for POST and PATCH requests.

|name
|string
a|The name of the SVM.

|nsswitch
|link:#nsswitch[nsswitch]
a|Name service switch configuration

|snapmirror
|link:#snapmirror[snapmirror]
a|Specifies attributes for SVM DR protection.

|snapshot_policy
|link:#snapshot_policy[snapshot_policy]
a|This is a reference to the Snapshot copy policy.

|state
|string
a|SVM State

```

```
|uuid
|string
a|The unique identifier of the SVM.
```

```
|volume_efficiency_policy
|link:#volume_efficiency_policy[volume_efficiency_policy]
a|This is a reference to the volume efficiency policy.
```

```
|===
```

```
[#job_link]
[.api-collapsible-fifth-title]
job_link
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|uuid
```

```
|string
```

```
a|The UUID of the asynchronous job that is triggered by a POST, PATCH, or
DELETE operation.
```

```
|===
```

```
[#error_arguments]
[.api-collapsible-fifth-title]
error_arguments
```

```
[cols=3*,options=header]
```

```
|===
```

```
|Name
```

```
|Type
```

```
|Description
```

```
|code
```

```
|string
```

```
a|Argument code
```

```

|message
|string
a|Message argument

|===

[#error]
[.api-collapsible-fifth-title]
error

[cols=3*,options=header]
|===
|Name
|Type
|Description

|arguments
|array[link:#error_arguments[error_arguments]]
a|Message arguments

|code
|string
a|Error code

|message
|string
a|Error message

|target
|string
a|The target parameter that caused the error.

|===

//end collapsible .Definitions block
====

:leveloffset: -1

```

:leveloffset: -1

= Security

:leveloffset: +1

[[IDbefd7f5cdfc71d4eeb4cb0dcdc045d24]]

= Security overview

== Overview

You can use ONTAP security APIs to manage security settings for the cluster and SVMs.

=== SAML

Configure the SAML 2.0 SP (Service Provider) protocol inside ONTAP. Doing so redirects the authentication task to a third-party Identity Provider (IDP) that can utilize any number of approaches for multi-factor authentication. After SAML authentication is enabled, all interactive web access (System Manager, SPI) is authenticated via SAML and a third-party IDP.

= Manage security-related operations

:leveloffset: +1

[[ID6a51a0a2d2d64f4df9ac752cd3a62894]]

= Security endpoint overview

:doctype: book

== Overview

You can use this API for various cluster-wide security-related operations.

== "onboard_key_manager_configurable_status" object

Use this API to retrieve details of whether or not the Onboard Key Manager can be configured on the cluster.

– GET /api/security

– GET

/api/security?fields=onboard_key_manager_configurable_status

== "software_data_encryption" object

Contains software data encryption related information.

A PATCH request on this API using the parameter

"software_data_encryption.conversion_enabled" triggers the conversion of all non-encrypted metadata volumes to encrypted metadata volumes and all non-NAE aggregates to NAE aggregates. For the conversion to start, the cluster must have either an Onboard or an external key manager set up and the aggregates should either be empty or have only metadata volumes. No data volumes should be present in any of the aggregates. For MetroCluster configurations, the PATCH request will fail if the cluster is in the switchover state.

The following API can be used to initiate software data encryption conversion.

– PATCH /api/security -d '{

"software_data_encryption.conversion_enabled" : true }'

== "fips" object

Contains FIPS mode information.

A PATCH request on this API using the parameter "fips.enabled" switches the system from using the default cryptographic module software implementations to validated ones or vice versa, where applicable. If the value of the parameter is "true" and unapproved algorithms are configured as permitted in relevant subsystems, those algorithms will be disabled in the relevant subsystem configurations. If "false", there will be no implied change to the relevant subsystem configurations.

– GET /api/security

– GET /api/security?fields=fips


```
&ndash; PATCH /api/security -d '{ "fips.enabled" : true }'
```

```
&ndash; PATCH /api/security -d '{ "fips.enabled" : false }'
```

== GET Examples

=== Retrieving information about the security configured on the cluster

The following example shows how to retrieve the configuration of the cluster.

The API:

GET /api/security:

The call:

```
curl -X GET 'https://<mgmt-ip>/api/security?fields=*' -H 'accept:
application/hal+json'
```

The response:

```
{
  "onboard_key_manager_configurable_status": {
    "supported": false,
    "message": "Onboard Key Manager cannot be configured on the cluster.
There are no self-encrypting disks in the cluster, and the following nodes
do not support volume granular encryption: ntap-vsim2.",
    "code": 65537300
  },
  "fips": {
    "enabled": false
  }
}
```

'''

PATCH Examples

Enabling software encryption conversion in the cluster

The following example shows how to convert all the aggregates and metadata volumes in the cluster from non-encrypted to encrypted.

The API:

PATCH /api/security

The call

```
curl -X PATCH "https://<mgmt_ip>/api/security" -d '{"software_data_encryption.conversion_enabled": true}'</mgmt_ip>
```

The response:

```
{
  "job": {
    "uuid": "ebcbd82d-1cd4-11ea-8f75-005056ac4adc",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/ebcbd82d-1cd4-11ea-8f75-005056ac4adc"
      }
    }
  }
}
```

This will return a job UUID. A subsequent GET for this job should return the details of the job.

The call

```
curl -X GET "https://<mgmt_ip>/api/cluster/jobs/ebcbd82d-1cd4-11ea-8f75-005056ac4adc"</mgmt_ip>
```

The response:

```
{
  "uuid": "ebcbd82d-1cd4-11ea-8f75-005056ac4adc",
  "description": "PATCH /api/security",
  "state": "success",
  "message": "success",
  "code": 0,
  "start_time": "2019-12-12T06:45:40-05:00",
  "end_time": "2019-12-12T06:45:40-05:00",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/ebcbd82d-1cd4-11ea-8f75-005056ac4adc"
    }
  }
}
```

Enabling FIPS mode in the cluster

The following example shows how to enable FIPS mode in the cluster.

The API:

PATCH /api/security

The call

```
curl -X PATCH "https://<mgmt_ip>/api/security" -d '{ "fips.enabled" : true }'</mgmt_ip>
```

The response:

```
{
  "job": {
    "uuid": "8e7f59ee-a9c4-4faa-9513-bef689bbf2c2",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/8e7f59ee-a9c4-4faa-9513-bef689bbf2c2"
      }
    }
  }
}
```

This will return a job UUID. A subsequent GET for this job UUID should return the details of the job.

The call

```
curl -X GET "https://<mgmt_ip>/api/cluster/jobs/8e7f59ee-a9c4-4faa-9513-bef689bbf2c2"</mgmt_ip>
```

The response:

```
{
  "uuid": "8e7f59ee-a9c4-4faa-9513-bef689bbf2c2",
  "description": "PATCH /api/security",
  "state": "success",
  "message": "success",
  "code": 0,
  "start_time": "2020-04-28T06:55:40-05:00",
  "end_time": "2020-04-28T06:55:41-05:00",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/8e7f59ee-a9c4-4faa-9513-bef689bbf2c2"
    }
  }
}
```

Retrieve information about security configured on the cluster

```
GET /security
```

Introduced In: 9.7

Retrieves information about the security configured on the cluster.

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
fips	fips	Cluster-wide Federal Information Processing Standards (FIPS) mode information.

Name	Type	Description
onboard_key_manager_configurable_status	onboard_key_manager_configurable_status	Indicates whether the Onboard Key Manager can be configured in the cluster.
software_data_encryption	software_data_encryption	Cluster-wide software data encryption related information.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "onboard_key_manager_configurable_status": {
    "code": "65537300",
    "message": "No platform support for volume encryption in following
nodes - node1, node2."
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

fips

Cluster-wide Federal Information Processing Standards (FIPS) mode information.

Name	Type	Description
enabled	boolean	Indicates whether or not the software FIPS mode is enabled on the cluster. Our FIPS compliance involves configuring the use of only approved algorithms in applicable contexts (for example TLS), as well as the use of formally validated cryptographic module software implementations, where applicable. The US government documents concerning FIPS 140-2 outline the relevant security policies in detail.

onboard_key_manager_configurable_status

Indicates whether the Onboard Key Manager can be configured in the cluster.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns a 0 if the Onboard Key Manager can be configured in the cluster.
message	string	Reason that Onboard Key Manager cannot be configured in the cluster.

Name	Type	Description
supported	boolean	Set to true if the Onboard Key Manager can be configured in the cluster.

software_data_encryption

Cluster-wide software data encryption related information.

Name	Type	Description
conversion_enabled	boolean	Indicates whether or not software encryption conversion is enabled on the cluster. A PATCH request initiates the conversion of all non-encrypted metadata volumes in the cluster to encrypted metadata volumes and all non-NAE aggregates to NAE aggregates. For the PATCH request to start, the cluster must have either an Onboard or an external key manager set up and the aggregates should either be empty or have only metadata volumes. No data volumes should be present in any of the aggregates in the cluster. For MetroCluster configurations, a PATCH request enables conversion on all the aggregates and metadata volumes of both local and remote clusters and is not allowed when the MetroCluster is in switchover state.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the software FIPS mode or enable conversion of non-encrypted metadata volumes non-NAE aggregates

PATCH /security

Introduced In: 9.8

Updates the software FIPS mode or enables conversion of non-encrypted metadata volumes to encrypted metadata volumes and non-NAE aggregates to NAE aggregates.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
fips	fips	Cluster-wide Federal Information Processing Standards (FIPS) mode information.
software_data_encryption	software_data_encryption	Cluster-wide software data encryption related information.

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
52428830	Cannot enable FIPS-compliant mode because the configured minimum security strength for certificates is not compatible.
52559974	Cannot enable FIPS-compliant mode because a certificate that is not FIPS-compliant is in use.
196608081	Cannot start software encryption conversion while there are data volumes in the cluster.
196608082	The operation is not valid when the MetroCluster is in switchover mode.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

fips

Cluster-wide Federal Information Processing Standards (FIPS) mode information.

Name	Type	Description
enabled	boolean	Indicates whether or not the software FIPS mode is enabled on the cluster. Our FIPS compliance involves configuring the use of only approved algorithms in applicable contexts (for example TLS), as well as the use of formally validated cryptographic module software implementations, where applicable. The US government documents concerning FIPS 140-2 outline the relevant security policies in detail.

onboard_key_manager_configurable_status

Indicates whether the Onboard Key Manager can be configured in the cluster.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns a 0 if the Onboard Key Manager can be configured in the cluster.
message	string	Reason that Onboard Key Manager cannot be configured in the cluster.
supported	boolean	Set to true if the Onboard Key Manager can be configured in the cluster.

software_data_encryption

Cluster-wide software data encryption related information.

Name	Type	Description
conversion_enabled	boolean	Indicates whether or not software encryption conversion is enabled on the cluster. A PATCH request initiates the conversion of all non-encrypted metadata volumes in the cluster to encrypted metadata volumes and all non-NAE aggregates to NAE aggregates. For the PATCH request to start, the cluster must have either an Onboard or an external key manager set up and the aggregates should either be empty or have only metadata volumes. No data volumes should be present in any of the aggregates in the cluster. For MetroCluster configurations, a PATCH request enables conversion on all the aggregates and metadata volumes of both local and remote clusters and is not allowed when the MetroCluster is in switchover state.

security_config

Name	Type	Description
fips	fips	Cluster-wide Federal Information Processing Standards (FIPS) mode information.
software_data_encryption	software_data_encryption	Cluster-wide software data encryption related information.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage security-related accounts

Security accounts endpoint overview

Overview

A valid user account is required to login to and provision, monitor, and manage the cluster. The scope of the management operation can be at the cluster level or at an individual SVM level. There is a need to create user accounts with specific privileges apart from the default user accounts, "admin", for cluster and "vsadmin" for SVM. Custom user accounts can be configured to perform specific (scoped) operations. User accounts can either be created locally (on the Netapp system) or referenced from an external directory server (NIS, LDAP, or Active Directory).

Apart from creation, modification, and deletion of a user account, locking and unlocking of a user account or resetting the password (for local accounts only) is possible.

A user account must be associated with the following before it can become operational:

1. A management application (SSH, HTTP, console, service_processor, and such like) for user login. HTTP enables REST API access.
2. Scope - either cluster or SVM.
3. Authentication source - password (local, NIS/LDAP, Active Directory), public/private key pair-based, certificate based.
4. RBAC role - determines what operations are permitted for the user account.

Restrictions

A number of internal/restricted account names, such as admin, diag, autosupport, and root cannot be used.

There must be at least one console cluster administrator account. Any attempt to delete the last remaining administrator account fails.

Multi-factor authentication is only possible for SSH application and the only combination possible is password (local or NIS/LDAP) and public key.

All authentication sources are not supported by all applications. You must select a compatible authentication method based on the application. The following types of authentications methods are supported:

Application	Supported Authentication Methods
console	password
service_processor	password
HTTP	password, domain, nsswitch, certificate
ONTAPI	password, domain, nsswitch, certificate
SSH	password, publickey (key pair), domain, nsswitch



In this table, "certificate" means security certificate, "domain" means that the user directory server is an external Active Directory, "nsswitch" means the directory server is an external NIS or LDAP server. At login time, the user is authenticated with these external directory servers which must be provisioned separately.

Examples

Creating a cluster-scoped user account

Specify the user account name, role name, and the tuples (of application and authentication methods) in the body of the POST request. The owner.uuid or owner.name are not required to be specified for a cluster-scoped user account.



Each entry in the applications array must be for a different application.

```
# The API:
POST "/api/security/accounts"

# The call to create a cluster user account with applications ssh, http
and password authentication scheme:
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"name":"cluster_user1","applications":[{"application":"ssh","authentication_methods":["password"],"second_authentication_method":"none"}, {"application":"http","authentication_methods":["password"]}], "role":"admin", "password":"p@ssw@rd123"}'
```

Note: The password is an optional parameter for creation and can be set later using a PATCH request. See the examples for modification of user account or password.

Creating an SVM-scoped user account

For an SVM-scoped account, specify either the SVM name as the owner.name or SVM uuid as the owner.uuid along with other parameters for the user account. These indicate the SVM for which the user account is being

created and can be obtained from the response body of GET performed on the `/api/svm/svms` API.

```
# The API:
POST "/api/security/accounts"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/accounts" -d
'{"owner":{"uuid":"aaef7c38-4bd3-11e9-b238-0050568e2e25"},"name":"svm_user1","applications":[{"application":"ssh","authentication_methods":["password"],"second_authentication_method":"none"}],"role":"vsadmin","password":"p@ssw@rd123"}'
```

Retrieving the configured user accounts

Use the following API to retrieve all of the user accounts or a filtered list of user accounts (by name, for a specific SVM, and so on).

```
# The API:
GET "/api/security/accounts"

# The call to retrieve all the user accounts configured in the cluster:
curl -X GET "https://<mgmt-ip>/api/security/accounts"

# The response:
{
  "records": [
    {
      "owner": {
        "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
        "name": "cluster1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
          }
        }
      },
      "name": "admin",
      "_links": {
        "self": {
          "href": "/api/security/accounts/2903de6f-4bd2-11e9-b238-0050568e2e25/admin"
        }
      }
    },
    {

```



```

    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "autosupport",
    "_links": {
      "self": {
        "href": "/api/security/accounts/2903de6f-4bd2-11e9-b238-0050568e2e25/autosupport"
      }
    }
  },
  {
    "owner": {
      "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
      "name": "cluster1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "cluster_user1",
    "_links": {
      "self": {
        "href": "/api/security/accounts/2903de6f-4bd2-11e9-b238-0050568e2e25/cluster_user1"
      }
    }
  },
  {
    "owner": {
      "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
        }
      }
    },
    "name": "svm_user1",

```

```

    "_links": {
      "self": {
        "href": "/api/security/accounts/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_user1"
      }
    },
    {
      "owner": {
        "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
        "name": "svml",
        "_links": {
          "self": {
            "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
          }
        }
      },
      "name": "vsadmin",
      "_links": {
        "self": {
          "href": "/api/security/accounts/aaef7c38-4bd3-11e9-b238-0050568e2e25/vsadmin"
        }
      }
    }
  ],
  "num_records": 5,
  "_links": {
    "self": {
      "href": "/api/security/accounts"
    }
  }
}

```

The scoped call to retrieve the configured cluster-scoped user accounts:
 curl -X GET "https://<mgmt-ip>/api/security/accounts/?scope=cluster"

The scoped call to retrieve the configured SVM-scoped user accounts:
 curl -X GET "https://<mgmt-ip>/api/security/accounts/?scope=svm"

The scoped call to retrieve the user accounts configured for the SVM "svml":
 curl -X GET "https://<mgmt-ip>/api/security/accounts/?owner.name=svml"

The scoped call to retrieve the user accounts configured with the "admin" role:

```
curl -X GET "https://<mgmt-ip>/api/security/accounts/?role=admin"
```

Retrieve user accounts in the cluster

GET /security/accounts

Introduced In: 9.6

Retrieves a list of user accounts in the cluster.

Related ONTAP commands

- `security login show`

Learn more

- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
scope	string	query	False	Filter by scope <ul style="list-style-type: none">• Introduced in: 9.7
owner.uuid	string	query	False	Filter by owner.uuid <ul style="list-style-type: none">• Introduced in: 9.7
owner.name	string	query	False	Filter by owner.name <ul style="list-style-type: none">• Introduced in: 9.7
locked	boolean	query	False	Filter by locked <ul style="list-style-type: none">• Introduced in: 9.7
name	string	query	False	Filter by name <ul style="list-style-type: none">• Introduced in: 9.7

Name	Type	In	Required	Description
comment	string	query	False	Filter by comment <ul style="list-style-type: none"> Introduced in: 9.7
applications.application	string	query	False	Filter by applications.application <ul style="list-style-type: none"> Introduced in: 9.7
applications.second_authentication_method	string	query	False	Filter by applications.second_authentication_method <ul style="list-style-type: none"> Introduced in: 9.7
applications.authentication_methods	string	query	False	Filter by applications.authentication_methods <ul style="list-style-type: none"> Introduced in: 9.7
role.name	string	query	False	Filter by role.name <ul style="list-style-type: none"> Introduced in: 9.7
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[account]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "applications": [
        {
          "application": "string",
          "authentication_methods": [
            "string"
          ],
          "second_authentication_method": "string"
        }
      ],
      "comment": "string",
      "name": "joe.smith",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "role": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "admin"
      }
    },
  ],
}
```

```
    "scope": "string"
  }
]
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

account_application

Name	Type	Description
application	string	Applications
authentication_methods	array[string]	
second_authentication_method	string	An optional additional authentication method for MFA. This only works with SSH as the application. It is ignored for all other applications.

owner

Owner name and UUID that uniquely identifies the user account.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	Role name

account

Name	Type	Description
_links	_links	
applications	array[account_application]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	owner	Owner name and UUID that uniquely identifies the user account.
role	role	
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Create a new user account

POST /security/accounts

Introduced In: 9.6

Creates a new user account.

Required parameters

- `name` - Account name to be created.
- `applications` - Array of one or more application tuples (of application and authentication methods).

Optional parameters

- `owner.name` or `owner.uuid` - Name or UUID of the SVM for an SVM-scoped user account. If not supplied, a cluster-scoped user account is created.
- `role` - RBAC role for the user account. Defaulted to `admin` for cluster user account and to `vsadmin` for SVM-scoped account.
- `password` - Password for the user account (if the authentication method is opted as password for one or more of applications).
- `second_authentication_method` - Needed for MFA and only supported for ssh application. Defaults to `none` if not supplied.
- `comment` - Comment for the user account (e.g purpose of this account).
- `locked` - Locks the account after creation. Defaults to `false` if not supplied.

Related ONTAP commands

- `security login create`

Learn more

- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
applications	array[account_application]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	owner	Owner name and UUID that uniquely identifies the user account.
password	string	Password for the account. The password can contain a mix of lower and upper case alphabetic characters, digits, and special characters.
role	role	
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

Example request

```
{
  "applications": [
    {
      "application": "string",
      "authentication_methods": [
        "string"
      ],
      "second_authentication_method": "string"
    }
  ],
  "comment": "string",
  "name": "joe.smith",
  "owner": {
    "name": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "password": "string",
  "role": {
    "name": "admin"
  },
  "scope": "string"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1261215	The role was not found.
1263343	Cannot lock user with password not set or non-password authentication method.
5636099	User creation with a non-admin role is not supported for service-processor application.

Error Code	Description
5636121	The user account name is reserved for use by the system.
5636126	Cannot create a user with the username or role as AutoSupport because it is reserved by the system.
5636140	Creating a login with application console for a data Vserver is not supported.
5636141	Creating a login with application service-processor for a data Vserver is not supported.
5636154	The second-authentication-method parameter is supported for ssh application.
5636155	The second-authentication-method parameter can be specified only if the authentication-method password or public key nsswitch.
5636156	The same value cannot be specified for the second-authentication-method and the authentication-method.
5636157	If the authentication-method is domain, the second-authentication-method cannot be specified.
5636164	If the value for either the authentication-method second-authentication-method is nsswitch or password, the other parameter must differ.
7077897	Invalid character in username.
7077898	The username must contain both letters and numbers.
7077899	The username does not meet length requirements.
7077906	A role with that name has not been defined for the Vserver.
7077918	The password cannot contain the username.
7077919	The minimum length for new password does not meet the policy.
7077920	A new password must have both letters and numbers.
7077921	The minimum number of special characters required do not meet the policy.
7077929	Cannot lock user with password not set or non-password authentication method.
7077940	The password exceeds the maximum supported length.
7077941	The defined password composition exceeds the maximum password length of 128 characters.
7078900	An admin password is not set. Set the password by including it in the request.

See Definitions

href

Name	Type	Description
href	string	

_links

account_application

Name	Type	Description
application	string	Applications
authentication_methods	array[string]	
second_authentication_method	string	An optional additional authentication method for MFA. This only works with SSH as the application. It is ignored for all other applications.

owner

Owner name and UUID that uniquely identifies the user account.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role

Name	Type	Description
name	string	Role name

account

Name	Type	Description
applications	array[account_application]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.

Name	Type	Description
name	string	User or group account name
owner	owner	Owner name and UUID that uniquely identifies the user account.
password	string	Password for the account. The password can contain a mix of lower and upper case alphabetic characters, digits, and special characters.
role	role	
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage scoped user accounts

Security accounts owner.uuid name endpoint overview

Overview

This API displays and manages the configuration of scoped user accounts.

Newly created user accounts might need to be updated for many reasons. For example, a user account might need to use a different application or its role might need to be modified. According to a policy, the password or authentication source of a user account might need to be changed, or a user account might need to be locked or deleted from the system. This API allows you to make these changes to user accounts.

Specify the owner UUID and the user account name in the URI path. The owner UUID corresponds to the UUID of the SVM for which the user account has been created and can be obtained from the response body of the GET request performed on one of the following APIs:

/api/security/accounts for all user accounts

/api/security/accounts/?scope=cluster for cluster-scoped user accounts

/api/security/accounts/?scope=svm for SVM-scoped accounts

/api/security/accounts/?owner.name={svm-name} for a specific SVM

This API response contains the complete URI for each user account that can be used.

Examples

Retrieving the user account details

```
# The API:
GET "/api/security/accounts/{owner.uuid}/{name}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/accounts/aef7c38-4bd3-11e9-b238-0050568e2e25/svm_user1"

# The response:
{
  "owner": {
    "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
      }
    }
  },
  "name": "svm_user1",
  "applications": [
    {
      "application": "ssh",
      "authentication_methods": [
        "password"
      ],
      "second_authentication_method": "none"
    }
  ]
}
```

```

],
"role": {
  "name": "vsadmin",
  "_links": {
    "self": {
      "href": "/api/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25/admin/roles/vsadmin"
    }
  }
},
"locked": false,
"scope": "svm",
"_links": {
  "self": {
    "href": "/api/security/accounts/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_user1"
  }
}
}

```

Updating the applications and role in a user account

Specify the desired configuration in the form of tuples (of applications and authentication methods) and the role. All other previously configured applications that are not specified in the "applications" parameter of the PATCH request will be de-provisioned for the user account.

```

# The API:
PATCH "/api/security/accounts/{owner.uuid}/{name}"

# The call to update the applications and role:
curl -X PATCH "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_user1" -d
'{"applications":[{"application":"http","authentication_methods":["domain"]}, {"application":"ontapi","authentication_methods":["password"]}], "role": {"name": "vsadmin-backup"}}'

# The call to update only the role:
curl -X PATCH "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_user1" -d '{"role":"vsadmin-protocol"}'

```

Updating the password for a user account

```
# The API:
PATCH "/api/security/accounts/{owner.uuid}/{name}"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_user1" -d '{"password":"newp@ssw@rd2"}'
```

Locking a user account

```
The API:
PATCH "/api/security/accounts/{owner.uuid}/{name}"

The call:
curl -X PATCH "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_user1" -d '{"locked":"true"}'
```

Deleting a user account

```
# The API:
DELETE "/api/security/accounts/{owner.uuid}/{name}"

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/accounts/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_user1"
```

Delete a user account

```
DELETE /security/accounts/{owner.uuid}/{name}
```

Introduced In: 9.6

Deletes a user account.

Required parameters

- `name` - Account name to be deleted.
- `owner.uuid` - UUID of the SVM housing the user account to be deleted.

Related ONTAP commands

- `security login delete`

Learn more

- [DOC /security/accounts/{owner.uuid}/{name}](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Account owner UUID
name	string	path	True	User account name

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5636098	Last unlocked account that has an admin role cannot be deleted.
5636125	The operation is not supported on system accounts.
5636146	Cannot delete the last console account with admin role.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a specific user account

GET /security/accounts/{owner.uuid}/{name}

Introduced In: 9.6

Retrieves a specific user account.

Related ONTAP commands

- `security login show`

Learn more

- [DOC /security/accounts/{owner.uuid}/{name}](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Account owner UUID
name	string	path	True	User account name
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
applications	array[account_application]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
name	string	User or group account name
owner	owner	Owner name and UUID that uniquely identifies the user account.
role	role	

Name	Type	Description
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "applications": [
    {
      "application": "string",
      "authentication_methods": [
        "string"
      ],
      "second_authentication_method": "string"
    }
  ],
  "comment": "string",
  "name": "joe.smith",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "role": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "admin"
  },
  "scope": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

account_application

Name	Type	Description
application	string	Applications
authentication_methods	array[string]	
second_authentication_method	string	An optional additional authentication method for MFA. This only works with SSH as the application. It is ignored for all other applications.

owner

Owner name and UUID that uniquely identifies the user account.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role

Name	Type	Description
_links	_links	
name	string	Role name

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a user account

PATCH /security/accounts/{owner.uuid}/{name}

Introduced In: 9.6

Updates a user account. Locks or unlocks a user account and/or updates the role, applications, and/or password for the user account.

Required parameters

- `name` - Account name to be updated.
- `owner.uuid` - UUID of the SVM housing the user account to be updated.

Optional parameters

- `applications` - Array of one or more tuples (of application and authentication methods).
- `role` - RBAC role for the user account.
- `password` - Password for the user account (if the authentication method is opted as password for one or more of applications).
- `second_authentication_method` - Needed for MFA and only supported for ssh application. Defaults to `none` if not supplied.
- `comment` - Comment for the user account (e.g purpose of this account).
- `locked` - Set to `true/false` to lock/unlock the account.

Related ONTAP commands

- `security login create`
- `security login modify`
- `security login password`
- `security login lock`
- `security login unlock`

Learn more

- [DOC /security/accounts/{owner.uuid}/{name}](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Account owner UUID
name	string	path	True	User account name

Request Body

Name	Type	Description
applications	array[account_application]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.
password	string	Password for the account. The password can contain a mix of lower and upper case alphabetic characters, digits, and special characters.
role	role	
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

Example request

```
{
  "applications": [
    {
      "application": "string",
      "authentication_methods": [
        "string"
      ],
      "second_authentication_method": "string"
    }
  ],
  "comment": "string",
  "password": "string",
  "role": {
    "name": "admin"
  },
  "scope": "string"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1261215	The role was not found.
1261218	The user was not found.
1263343	Cannot lock user with password not set or non-password authentication method.
5636096	Cannot perform the operation for this user account since the password is not set.
5636097	The operation for user account failed since user password is not set.
5636100	User modification is not supported for service-processor application.

Error Code	Description
5636125	The operation not supported on AutoSupport user account which is reserved.
5636129	The role does not exist.
5636154	The second-authentication-method parameter is supported for ssh application.
5636155	The second-authentication-method parameter can be specified only if the authentication-method password or public key nsswitch.
5636156	Same value cannot be specified for the second-authentication-method and the authentication-method.
5636157	If the authentication-method is domain, the second-authentication-method cannot be specified.
5636159	For a given user and application, if the second-authentication-method is specified, only one such login entry is supported.
5636164	If the value for either the authentication-method second-authentication-method is nsswitch or password, the other parameter must differ.
5636174	You are not authorized to change the password for other users.
7077896	Cannot lock the account of the last console admin user.
7077906	A role with that name has not been defined for the Vserver.
7077911	The user is not configured to use the password authentication method.
7077918	The password cannot contain the username.
7077919	The minimum length for new password does not meet the policy.
7077920	The new password must have both letters and numbers.
7077921	The minimum number of special characters required do not meet the policy.
7077924	The new password must be different than last N passwords.
7077925	The new password must be different to the old password.
7077929	Cannot lock user with password not set or non-password authentication method.
7077940	The password exceeds maximum supported length.

Error Code	Description
7077941	Defined password composition exceeds the maximum password length of 128 characters.
7078900	An aAdmin password is not set. Set the password by including it in the request.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

account_application

Name	Type	Description
application	string	Applications
authentication_methods	array[string]	
second_authentication_method	string	An optional additional authentication method for MFA. This only works with SSH as the application. It is ignored for all other applications.

owner

Owner name and UUID that uniquely identifies the user account.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role

Name	Type	Description
name	string	Role name

account

Name	Type	Description
applications	array[account_application]	
comment	string	Optional comment for the user account.
locked	boolean	Locked status of the account.

Name	Type	Description
password	string	Password for the account. The password can contain a mix of lower and upper case alphabetic characters, digits, and special characters.
role	role	
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View and update audit settings

Security audit endpoint overview

Overview

This API controls what is logged to the audit log files. All operations that make changes are always logged and cannot be disabled. The PATCH request updates administrative audit settings for GET requests. All fields are optional for a PATCH request. A GET request retrieves administrative audit settings for GET requests.

+

Examples

Retrieving administrative audit settings for GET requests

The following example shows the administrative audit settings for GET requests.

+

```
# The API:
/api/security/audit

# The call:
curl -X GET "https://<cluster-ip>/api/security/audit"

# The response:
{
  "cli": false,
  "http": false,
  "ontapi": false,
  "_links": {
    "self": {
      "href": "/api/security/audit"
    }
  }
}
```

Updating administrative audit settings for GET requests

The following example updates the administrative audit settings for GET requests

+

```
# The API:
/api/security/audit

# The call:
curl -X PATCH "https://<cluster-ip>/api/security/audit" -d
'{"cli": "false", "http": "true", "ontapi": "true"}'
```

Retrieve administrative audit settings for GET requests

GET /security/audit

Introduced In: 9.6

Retrieves administrative audit settings for GET requests.

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
cli	boolean	Enable auditing of CLI GET Operations. Valid in PATCH
http	boolean	Enable auditing of HTTP GET Operations. Valid in PATCH
ontapi	boolean	Enable auditing of ONTAP API GET operations. Valid in PATCH <ul style="list-style-type: none">• Introduced in: 9.6

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update administrative audit settings for GET requests

PATCH /security/audit

Introduced In: 9.6

Updates administrative audit settings for GET requests.
All of the fields are optional. An empty body will make no changes.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
cli	boolean	Enable auditing of CLI GET Operations. Valid in PATCH
http	boolean	Enable auditing of HTTP GET Operations. Valid in PATCH
ontapi	boolean	<p>Enable auditing of ONTAP API GET operations. Valid in PATCH</p> <ul style="list-style-type: none"> • Introduced in: 9.6

Response

Status: 202, Accepted

Name	Type	Description
cli	boolean	Enable auditing of CLI GET Operations. Valid in PATCH
http	boolean	Enable auditing of HTTP GET Operations. Valid in PATCH
ontapi	boolean	Enable auditing of ONTAP API GET operations. Valid in PATCH <ul style="list-style-type: none">• Introduced in: 9.6

Error

Status: Default, Error

Definitions

See Definitions

security_audit

Name	Type	Description
cli	boolean	Enable auditing of CLI GET Operations. Valid in PATCH
http	boolean	Enable auditing of HTTP GET Operations. Valid in PATCH
ontapi	boolean	Enable auditing of ONTAP API GET operations. Valid in PATCH <ul style="list-style-type: none">Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Forward audit logs to syslog/splunk servers

Security audit destinations endpoint overview

Overview

This API controls the forwarding of audit log information to remote syslog/splunk servers. Multiple destinations can be configured and all audit records are forwarded to all destinations.

A GET operation retrieves information about remote syslog/splunk server destinations.

A POST operation creates a remote syslog/splunk server destination.

A GET operation on `/security/audit/destinations/{address}/{port}` retrieves information about the syslog/splunk server destination given its address and port number.

A PATCH operation on `/security/audit/destinations/{address}/{port}` updates information about the syslog/splunk server destination given its address and port number.

A DELETE operation on `/security/audit/destinations/{address}/{port}` deletes a syslog/splunk server destination given its address and port number.

Overview of fields used for creating a remote syslog/splunk destination

The fields used for creating a remote syslog/splunk destination fall into the following categories

Required properties

All of the following fields are required for creating a remote syslog/splunk destination

- `address`

Optional properties

All of the following fields are optional for creating a remote syslog/splunk destination

- `port`
- `protocol`
- `facility`
- `verify_server`
- `+`

Examples

Retrieving remote syslog/splunk server destinations

The following example shows remote syslog/splunk server destinations

+


```
# The API:
/api/security/audit/destinations

# The call:
curl -X GET "https://<cluster-ip>/api/security/audit/destinations"

# The response:
{
  "records": [
    {
      "address": "1.1.1.1",
      "port": 514,
      "_links": {
        "self": {
          "href": "/api/security/audit/destinations/1.1.1.1/514"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/audit/destinations"
    }
  }
}
```

Creating remote syslog/splunk server destinations

The following example creates remote syslog/splunk server destinations.

+

```
# The API:
/api/security/audit/destinations

# The call:
curl -X POST "https://<cluster-
ip>/api/security/audit/destinations?force=true" -d '{ "address":
"1.1.1.1", "port": 514, "protocol": "udp_unencrypted", "facility":
"kern"}'
```

Retrieving a remote syslog/splunk server destination given its destination address and port number

The following example retrieves a remote syslog/splunk server destination given its destination address and port number.

+

```
# The API:
/api/security/audit/destinations/{address}/{port}

# The call:
curl -X GET "https://<cluster-
ip>/api/security/audit/destinations/1.1.1.1/514"

# The response:
{
  "address": "1.1.1.1",
  "port": 514,
  "protocol": "udp_unencrypted",
  "facility": "kern",
  "verify_server": false,
  "_links": {
    "self": {
      "href": "/api/security/audit/destinations/1.1.1.1/514"
    }
  }
}
```

Updating a remote syslog/splunk server destination given its destination address and port number

The following example updates a remote syslog/splunk server destination configuration given its destination address and port number.

+

```
# The API:
/api/security/audit/destinations/{address}/{port}

# The call:
curl -X PATCH "https://<cluster-
ip>/api/security/audit/destinations/1.1.1.1/514" -d '{"facility":
"user}'
```

Deleting a remote syslog/splunk server destination given its destination address and port number

The following example deletes a remote syslog/splunk server destination configuration given its destination address and port number.

+

```
# The API:
/api/security/audit/destinations/{address}/{port}

# The call:
curl -X DELETE "https://<cluster-
ip>/api/security/audit/destinations/1.1.1.1/514"
```

Define a remote syslog or splunk server to receive audit information

GET /security/audit/destinations

Introduced In: 9.6

Defines a remote syslog/splunk server for sending audit information to.

Parameters

Name	Type	In	Required	Description
facility	string	query	False	Filter by facility
verify_server	boolean	query	False	Filter by verify_server
protocol	string	query	False	Filter by protocol
port	integer	query	False	Filter by port
address	string	query	False	Filter by address
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Name	Type	In	Required	Description
Specify the fields to return.	max_records	integer	query	False
Limit the number of records returned.	return_timeout	integer	query	False
<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0 	return_records	boolean	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[security_audit_log_forward]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "address": "string",
      "facility": "string",
      "protocol": "string"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

security_audit_log_forward

Name	Type	Description
address	string	Destination syslog splunk host to forward audit records to. This can be an IP address (IPv4 IPv6) or a hostname.
facility	string	This is the standard Syslog Facility value that is used when sending audit records to a remote server.
port	integer	Destination Port. The default port depends on the protocol chosen: For un-encrypted destinations the default port is 514. For encrypted destinations the default port is 6514.
protocol	string	Log forwarding protocol
verify_server	boolean	This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Define remote syslog or splunk server information

POST /security/audit/destinations

Introduced In: 9.6

Configures remote syslog/splunk server information.

Required properties

All of the following fields are required for creating a remote syslog/splunk destination

- address

Optional properties

All of the following fields are optional for creating a remote syslog/splunk destination

- port
- protocol
- facility
- verify_server (Can only be "true" when protocol is "tcp_encrypted")

Parameters

Name	Type	In	Required	Description
force	boolean	query	False	<p>Skip the Connectivity Test</p> <ul style="list-style-type: none"> • Default value:
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
address	string	Destination syslog splunk host to forward audit records to. This can be an IP address (IPv4 IPv6) or a hostname.
facility	string	This is the standard Syslog Facility value that is used when sending audit records to a remote server.
port	integer	Destination Port. The default port depends on the protocol chosen: For un-encrypted destinations the default port is 514. For encrypted destinations the default port is 6514.
protocol	string	Log forwarding protocol
verify_server	boolean	This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate.

Example request

```
{
  "address": "string",
  "facility": "string",
  "protocol": "string"
}
```

Response

```
Status: 202, Accepted
```

Name	Type	Description
num_records	integer	Number of records
records	array[security_audit_log_forward]	

Example response

```
{
  "records": [
    {
      "address": "string",
      "facility": "string",
      "protocol": "string"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
15661	The object specified could not be found
13114	Internal error
13115	Invalid input
4522285	Server verification cannot be enabled because it requires a protocol with encryption. Encryption can be selected using the protocol field.
9240603	Cannot ping destination host. Verify connectivity to desired host or skip the connectivity check with the -force parameter.
327698	Failed to create RPC client to destination host
9240609	Cannot connect to destination host.
9240604	Cannot resolve the destination host.

Definitions

See Definitions

security_audit_log_forward

Name	Type	Description
address	string	Destination syslog splunk host to forward audit records to. This can be an IP address (IPv4 IPv6) or a hostname.
facility	string	This is the standard Syslog Facility value that is used when sending audit records to a remote server.
port	integer	Destination Port. The default port depends on the protocol chosen: For un-encrypted destinations the default port is 514. For encrypted destinations the default port is 6514.
protocol	string	Log forwarding protocol
verify_server	boolean	This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate.

href

Name	Type	Description
href	string	

_links

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete remote syslog or splunk server information

DELETE /security/audit/destinations/{address}/{port}

Introduced In: 9.6

Deletes remote syslog/splunk server information.

Parameters

Name	Type	In	Required	Description
address	string	path	True	IP address of remote syslog/splunk server.
port	integer	path	True	Port number of remote syslog/splunk server.

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve remote syslog or splunk server information

GET /security/audit/destinations/{address}/{port}

Introduced In: 9.6

Defines a remote syslog/splunk server for sending audit information to.

Parameters

Name	Type	In	Required	Description
address	string	path	True	IP address of remote syslog/splunk server.
port	integer	path	True	Port number of remote syslog/splunk server.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
address	string	Destination syslog splunk host to forward audit records to. This can be an IP address (IPv4 IPv6) or a hostname.
facility	string	This is the standard Syslog Facility value that is used when sending audit records to a remote server.
port	integer	Destination Port. The default port depends on the protocol chosen: For un-encrypted destinations the default port is 514. For encrypted destinations the default port is 6514.
protocol	string	Log forwarding protocol

Name	Type	Description
verify_server	boolean	This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate.

Example response

```
{
  "address": "string",
  "facility": "string",
  "protocol": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update remote syslog or splunk server information

PATCH /security/audit/destinations/{address}/{port}

Introduced In: 9.6

Updates remote syslog/splunk server information.

Parameters

Name	Type	In	Required	Description
address	string	path	True	IP address of remote syslog/splunk server.
port	integer	path	True	Port number of remote syslog/splunk server.

Request Body

Name	Type	Description
facility	string	This is the standard Syslog Facility value that is used when sending audit records to a remote server.
verify_server	boolean	This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate.

Example request

```
{  
  "facility": "string"  
}
```

Response

Status: 200, Ok

Name	Type	Description
facility	string	This is the standard Syslog Facility value that is used when sending audit records to a remote server.
verify_server	boolean	This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate.

Example response

```
{
  "facility": "string"
}
```

Error

Status: Default, Default

Definitions

See Definitions

security_audit_log_forward

Name	Type	Description
facility	string	This is the standard Syslog Facility value that is used when sending audit records to a remote server.
verify_server	boolean	This is only applicable when the protocol is tcp_encrypted. This controls whether the remote server's certificate is validated. Setting "verify_server" to "true" will enforce validation of remote server's certificate. Setting "verify_server" to "false" will not enforce validation of remote server's certificate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View administrative audit logs

Security audit messages endpoint overview

Overview

These APIs return audit log records. The GET requests retrieves all audit log records. An audit log record contains information such as timestamp, node name, index and so on.

+

Example

Retrieving audit log records

The following example shows the audit log records.

+

```
# The API:
/api/security/audit/messages

# The call:
curl -X GET "https://<cluster-ip>/api/security/audit/messages"

# The response:
{
  "records": [
    {
      "timestamp": "2019-03-08T11:03:32-05:00",
      "node": {
        "name": "node1",
        "uuid": "bc9af9da-41bb-11e9-a3db-005056bb27cf",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/bc9af9da-41bb-11e9-a3db-005056bb27cf"
          }
        }
      },
      "index": 4294967299,
      "application": "http",
      "location": "172.21.16.89",
      "user": "admin",
      "input": "GET /api/security/audit/destinations/",
      "state": "pending",
      "scope": "cluster"
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/audit/messages"
    }
  }
}
```

Retrieve the administrative audit log viewer

GET /security/audit/messages

Introduced In: 9.6

Retrieves the administrative audit log viewer.

Parameters

Name	Type	In	Required	Description
command_id	string	query	False	Filter by command_id
application	string	query	False	Filter by application
svm.name	string	query	False	Filter by svm.name
input	string	query	False	Filter by input
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
index	integer	query	False	Filter by index
user	string	query	False	Filter by user
session_id	string	query	False	Filter by session_id
state	string	query	False	Filter by state
message	string	query	False	Filter by message
scope	string	query	False	Filter by scope
location	string	query	False	Filter by location
timestamp	string	query	False	Filter by timestamp
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[security_audit_log]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "application": "string",
      "command_id": "string",
      "index": 0,
      "input": "string",
      "location": "string",
      "message": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
      },
      "scope": "string",
      "session_id": "string",
      "state": "string",
      "svm": {
        "name": "string"
      },
      "timestamp": "string",
      "user": "string"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

node

Node where the audit message resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

This is the SVM through which the user connected.

Name	Type	Description
name	string	

security_audit_log

Name	Type	Description
_links	_links	
application	string	This identifies the "application" by which the request was processed.

Name	Type	Description
command_id	string	This is the command ID for this request. Each command received on a CLI session is assigned a command ID. This enables you to correlate a request and response.
index	integer	Internal index for accessing records with same time/node. This is a 64 bit unsigned value.
input	string	The request.
location	string	This identifies the location of the remote user. This is an IP address or "console".
message	string	This is an optional field that might contain "error" or "additional information" about the status of a command.
node	node	Node where the audit message resides.
scope	string	Set to "svm" when the request is on a data SVM; otherwise set to "cluster".
session_id	string	This is the session ID on which the request is received. Each SSH session is assigned a session ID. Each http/ontapi/snmp request is assigned a unique session ID.
state	string	State of of this request.
svm	svm	This is the SVM through which the user connected.
timestamp	string	Log entry timestamp. Valid in URL
user	string	Username of the remote user.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage data SVM account information

Security authentication cluster ad-proxy endpoint overview

Overview

This API configures data SVM account information at the Active Directory. For Active Directory domain-based authentication for cluster accounts, a data SVM must be configured and registered as a machine account at the Active Directory. All authentication requests are proxied through this SVM.

Examples

Creating a data SVM proxy for domain-based authentication for cluster accounts

```
# The API:
POST  "/api/security/authentication/cluster/ad-proxy"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/ad-proxy" -d '{"svm.uuid":"13f87d78-70c7-11e9-b722-0050568ec89f"}'
```

Updating a data SVM proxy for domain-based authentication for cluster accounts

```
# The API:
PATCH "/api/security/authentication/cluster/ad-proxy"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/ad-proxy" -d '{"svm.uuid":"13f87d78-70c7-11e9-b722-0050568ec89f"}'
```

Retrieving a data SVM proxy for domain-based authentication for cluster accounts

```
# The API:
GET "/api/security/authentication/cluster/ad-proxy"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/ad-proxy"

# The response:
{
  "svm": {
    "uuid": "512eab7a-6bf9-11e9-a896-005056bb9ce1",
    "name": "vs2",
    "_links": {
      "self": {
        "href": "/api/svm/svms/512eab7a-6bf9-11e9-a896-005056bb9ce1"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/security/authentication/cluster/ad-proxy"
    }
  }
}
```

Delete a data SVM configured as a tunnel

DELETE /security/authentication/cluster/ad-proxy

Introduced In: 9.7

Deletes the data SVM configured as a tunnel for Active Directory based authentication for cluster user accounts.

Related ONTAP commands

- `security login domain-tunnel delete`

Learn more

- [DOC /security/authentication/cluster/ad-proxy](#)
- [DOC /security/accounts](#)

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve SVM information configured as an Active Directory domain-tunnel

GET /security/authentication/cluster/ad-proxy

Introduced In: 9.7

Retrieves SVM information configured as an Active Directory domain-tunnel.

Related ONTAP commands

- `security login domain-tunnel show`

Learn more

- [DOC /security/authentication/cluster/ad-proxy](#)
- [DOC /security/accounts](#)

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
svm	svm	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a data SVM configured as a tunnel

PATCH /security/authentication/cluster/ad-proxy

Introduced In: 9.7

Updates the data SVM configured as a tunnel for Active Directory based authentication for cluster user accounts.

Related ONTAP commands

- `security login domain-tunnel modify`

Learn more

- [DOC /security/authentication/cluster/ad-proxy](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
svm	svm	

Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cluster_ad_proxy

The SVM configured as proxy for Active Directory authentication of cluster accounts.

Name	Type	Description
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Configure a data SVM as a proxy

POST /security/authentication/cluster/ad-proxy

Introduced In: 9.7

Configures a data SVM as a proxy for Active Directory based authentication for cluster user accounts.

Required properties

- svm.name or svm.uuid - Name and UUID of the SVM for a cluster user account.

Related ONTAP commands

- security login domain-tunnel create

Learn more

- [DOC /security/authentication/cluster/ad-proxy](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
svm	svm	

Example request

```
{
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 201, Created

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

cluster_ad_proxy

The SVM configured as proxy for Active Directory authentication of cluster accounts.

Name	Type	Description
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage LDAP server configuration

Security authentication cluster LDAP endpoint overview

Overview

LDAP servers are used to centrally maintain user information. LDAP configurations must be set up to look up information stored in the LDAP directory on the external LDAP servers. This API is used to retrieve and manage cluster LDAP server configurations.

Examples

Retrieving the cluster LDAP information

The cluster LDAP GET request retrieves the LDAP configuration of the cluster.

The following example shows how a GET request is used to retrieve the cluster LDAP information:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/ldap"
-H "accept: application/hal+json"

# The response:
{
  "servers": [
    "10.10.10.10",
    "domainB.example.com"
  ],
  "schema": "ad_idmu",
  "port": 389,
  "min_bind_level": "anonymous",
  "bind_dn": "cn=Administrators,cn=users,dc=domainA,dc=example,dc=com",
  "base_dn": "dc=domainA,dc=example,dc=com",
  "base_scope": "subtree",
  "use_start_tls": true,
  "session_security": "none",
  "_links": {
    "self": {
      "href": "/api/security/authentication/cluster/ldap"
    }
  }
}
```

Creating the cluster LDAP configuration

The cluster LDAP POST operation creates an LDAP configuration for the cluster.

The following example shows how to issue a POST request with all of the fields specified:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/ldap"
-H "accept: application/hal+json" -H "Content-Type: application/json" -d
"{ \"servers\": [ \"10.10.10.10\", \"domainB.example.com\" ], \"schema\":
\"ad_idmu\", \"port\": 389, \"min_bind_level\": \"anonymous\",
\"bind_dn\": \"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\",
\"base_scope\": \"subtree\", \"use_start_tls\": false,
\"session_security\": \"none\"}"
```

The following example shows how to issue a POST request with a number of optional fields not specified:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/ldap"
-H "accept: application/hal+json" -H "Content-Type: application/json" -d
"{ \"port\": 389, \"bind_dn\":
\"cn=Administrators,cn=users,dc=domainA,dc=example,dc=com\",
\"bind_password\": \"abc\", \"base_dn\": \"dc=domainA,dc=example,dc=com\",
\"session_security\": \"none\"}"
```

Updating the cluster LDAP configuration

The cluster LDAP PATCH request updates the LDAP configuration of the cluster.

The following example shows how a PATCH request is used to update the cluster LDAP configuration:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/ldap"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
\"servers\": [ \"55.55.55.55\" ], \"schema\": \"ad_idmu\", \"port\": 636,
\"use_start_tls\": false }"
```

Deleting the cluster LDAP configuration

The cluster LDAP DELETE request deletes the LDAP configuration of the cluster.

The following example shows how a DELETE request is used to delete the cluster LDAP configuration:

```
# The API:
/api/security/authentication/cluster/ldap

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/security/authentication/cluster/ldap" -H "accept:
application/hal+json"
```

Delete the LDAP configuration for the cluster

DELETE /security/authentication/cluster/ldap

Introduced In: 9.6

Deletes the LDAP configuration of the cluster.

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the LDAP configuration for the cluster

GET /security/authentication/cluster/ldap

Introduced In: 9.6

Retrieves the cluster LDAP configuration.

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
base_dn	string	Specifies the default base DN for all searches.
base_scope	string	<p>Specifies the default search scope for LDAP queries:</p> <ul style="list-style-type: none"> • base - search the named entry only • onelevel - search all entries immediately below the DN • subtree - search the named DN entry and the entire subtree below the DN
bind_dn	string	Specifies the user that binds to the LDAP servers.
min_bind_level	string	<p>The minimum bind authentication level. Possible values are:</p> <ul style="list-style-type: none"> • anonymous - anonymous bind • simple - simple bind • sasl - Simple Authentication and Security Layer (SASL) bind
port	integer	The port used to connect to the LDAP Servers.
schema	string	<p>The name of the schema template used by the SVM.</p> <ul style="list-style-type: none"> • AD-IDMU - Active Directory Identity Management for UNIX • AD-SFU - Active Directory Services for UNIX • MS-AD-BIS - Active Directory Identity Management for UNIX • RFC-2307 - Schema based on RFC 2307 • Custom schema
servers	array[string]	

Name	Type	Description
session_security	string	Specifies the level of security to be used for LDAP communications: <ul style="list-style-type: none"> • none - no signing or sealing • sign - sign LDAP traffic • seal - seal and sign LDAP traffic
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "min_bind_level": "string",
  "port": "389",
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the LDAP configuration for the cluster

PATCH /security/authentication/cluster/ldap

Introduced In: 9.6

Both mandatory and optional parameters of the LDAP configuration can be updated.

IPv6 must be enabled if IPv6 family addresses are specified. Configuring more than one LDAP server is recommended to avoid a single point of failure. Both FQDNs and IP addresses are supported for the `servers` property.

The LDAP servers are validated as part of this operation. LDAP validation fails in the following scenarios:

1. The server does not have LDAP installed.
2. The server is invalid.

3. The server is unreachable.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
base_dn	string	Specifies the default base DN for all searches.
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none">• base - search the named entry only• onelevel - search all entries immediately below the DN• subtree - search the named DN entry and the entire subtree below the DN
bind_dn	string	Specifies the user that binds to the LDAP servers.
bind_password	string	Specifies the bind password for the LDAP servers.
min_bind_level	string	The minimum bind authentication level. Possible values are: <ul style="list-style-type: none">• anonymous - anonymous bind• simple - simple bind• sasl - Simple Authentication and Security Layer (SASL) bind
port	integer	The port used to connect to the LDAP Servers.

Name	Type	Description
schema	string	<p>The name of the schema template used by the SVM.</p> <ul style="list-style-type: none"> • AD-IDMU - Active Directory Identity Management for UNIX • AD-SFU - Active Directory Services for UNIX • MS-AD-BIS - Active Directory Identity Management for UNIX • RFC-2307 - Schema based on RFC 2307 • Custom schema
servers	array[string]	
session_security	string	<p>Specifies the level of security to be used for LDAP communications:</p> <ul style="list-style-type: none"> • none - no signing or sealing • sign - sign LDAP traffic • seal - seal and sign LDAP traffic
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

Example request

```
{
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "bind_password": "string",
  "min_bind_level": "string",
  "port": "389",
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
4915203	The specified LDAP schema does not exist.
4915208	The specified LDAP servers contain duplicate server entries.
4915229	DNS resolution failed due to an internal error. Contact technical support if this issue persists.
4915231	DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured.
23724132	DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured.
4915234	Specified LDAP server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast.
4915248	LDAP servers cannot be empty or "-". Specified FQDN is not valid because it is empty or "-" or it contains either special characters or "-" at the start or end of the domain.
4915251	STARTTLS and LDAPS cannot be used together
4915257	The LDAP configuration is not valid. Verify that the Distinguished Names and bind password are correct.
4915258	The LDAP configuration is not valid. Verify that the servers are reachable and that the network configuration is correct.
23724130	Cannot use an IPv6 name server address because there are no IPv6 interfaces.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster_ldap

Name	Type	Description
base_dn	string	Specifies the default base DN for all searches.
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none">• base - search the named entry only• onelevel - search all entries immediately below the DN• subtree - search the named DN entry and the entire subtree below the DN
bind_dn	string	Specifies the user that binds to the LDAP servers.
bind_password	string	Specifies the bind password for the LDAP servers.
min_bind_level	string	The minimum bind authentication level. Possible values are: <ul style="list-style-type: none">• anonymous - anonymous bind• simple - simple bind• sasl - Simple Authentication and Security Layer (SASL) bind
port	integer	The port used to connect to the LDAP Servers.

Name	Type	Description
schema	string	<p>The name of the schema template used by the SVM.</p> <ul style="list-style-type: none"> • AD-IDMU - Active Directory Identity Management for UNIX • AD-SFU - Active Directory Services for UNIX • MS-AD-BIS - Active Directory Identity Management for UNIX • RFC-2307 - Schema based on RFC 2307 • Custom schema
servers	array[string]	
session_security	string	<p>Specifies the level of security to be used for LDAP communications:</p> <ul style="list-style-type: none"> • none - no signing or sealing • sign - sign LDAP traffic • seal - seal and sign LDAP traffic
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Create the LDAP configuration for the cluster

POST /security/authentication/cluster/ldap

Introduced In: 9.6

A cluster can have only one LDAP configuration. IPv6 must be enabled if IPv6 family addresses are specified.

Required properties

- `servers` - List of LDAP servers used for this client configuration.
- `bind_dn` - Specifies the user that binds to the LDAP servers.
- `base_dn` - Specifies the default base DN for all searches.

Recommended optional properties

- `schema` - Schema template name.
- `port` - Port used to connect to the LDAP Servers.
- `min_bind_level` - Minimum bind authentication level.
- `bind_password` - Specifies the bind password for the LDAP servers.
- `base_scope` - Specifies the default search scope for LDAP queries.
- `use_start_tls` - Specifies whether or not to use Start TLS over LDAP connections.
- `session_security` - Specifies the level of security to be used for LDAP communications.

Default property values

- `schema` - *RFC-2307*
- `port` - *389*
- `min_bind_level` - *simple*
- `base_scope` - *subtree*
- `use_start_tls` - *false*
- `session_security` - *none*

Configuring more than one LDAP server is recommended to avoid a single point of failure. Both FQDNs and IP addresses are supported for the `servers` property.

The LDAP servers are validated as part of this operation. LDAP validation fails in the following scenarios:

1. The server does not have LDAP installed.
2. The server is invalid.
3. The server is unreachable.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
base_dn	string	Specifies the default base DN for all searches.
base_scope	string	<p>Specifies the default search scope for LDAP queries:</p> <ul style="list-style-type: none">• base - search the named entry only• onelevel - search all entries immediately below the DN• subtree - search the named DN entry and the entire subtree below the DN
bind_dn	string	Specifies the user that binds to the LDAP servers.
bind_password	string	Specifies the bind password for the LDAP servers.
min_bind_level	string	<p>The minimum bind authentication level. Possible values are:</p> <ul style="list-style-type: none">• anonymous - anonymous bind• simple - simple bind• sasl - Simple Authentication and Security Layer (SASL) bind

Name	Type	Description
port	integer	The port used to connect to the LDAP Servers.
schema	string	<p>The name of the schema template used by the SVM.</p> <ul style="list-style-type: none"> • AD-IDMU - Active Directory Identity Management for UNIX • AD-SFU - Active Directory Services for UNIX • MS-AD-BIS - Active Directory Identity Management for UNIX • RFC-2307 - Schema based on RFC 2307 • Custom schema
servers	array[string]	
session_security	string	<p>Specifies the level of security to be used for LDAP communications:</p> <ul style="list-style-type: none"> • none - no signing or sealing • sign - sign LDAP traffic • seal - seal and sign LDAP traffic
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

Example request

```
{
  "base_dn": "string",
  "base_scope": "string",
  "bind_dn": "string",
  "bind_password": "string",
  "min_bind_level": "string",
  "port": "389",
  "schema": "string",
  "servers": [
    "string"
  ],
  "session_security": "string"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of LDAP records.
records	array[ldap_service]	

Example response

```
{
  "records": [
    {
      "ad_domain": "string",
      "base_dn": "string",
      "base_scope": "string",
      "bind_dn": "string",
      "bind_password": "string",
      "min_bind_level": "string",
      "port": "389",
      "preferred_ad_servers": [
        "string"
      ],
      "schema": "string",
      "servers": [
        "string"
      ],
      "session_security": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
4915203	The specified LDAP schema does not exist.
4915207	The specified LDAP servers contain duplicate server entries.
4915229	DNS resolution failed due to an internal error. Contact technical support if this issue persists.

Error Code	Description
4915231	DNS resolution failed for one or more of the specified LDAP servers. Verify that a valid DNS server is configured.
23724132	DNS resolution failed for all the specified LDAP servers. Verify that a valid DNS server is configured.
4915234	The specified LDAP server is not supported because it is one of the following: multicast, loopback, 0.0.0.0, or broadcast.
4915248	LDAP servers cannot be empty or "-". Specified FQDN is invalid because it is empty or "-" or it contains either special characters or "-" at the start or end of the domain.
4915251	STARTTLS and LDAPS cannot be used together.
4915257	The LDAP configuration is invalid. Verify that bind-dn and bind password are correct.
4915258	The LDAP configuration is invalid. Verify that the servers are reachable and that the network configuration is correct.
13434916	The SVM is in the process of being created. Wait a few minutes, and then try the command again.
23724130	Cannot use an IPv6 name server address because there are no IPv6 interfaces.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster_ldap

Name	Type	Description
base_dn	string	Specifies the default base DN for all searches.
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none">• base - search the named entry only• onelevel - search all entries immediately below the DN• subtree - search the named DN entry and the entire subtree below the DN
bind_dn	string	Specifies the user that binds to the LDAP servers.
bind_password	string	Specifies the bind password for the LDAP servers.
min_bind_level	string	The minimum bind authentication level. Possible values are: <ul style="list-style-type: none">• anonymous - anonymous bind• simple - simple bind• sasl - Simple Authentication and Security Layer (SASL) bind
port	integer	The port used to connect to the LDAP Servers.

Name	Type	Description
schema	string	<p>The name of the schema template used by the SVM.</p> <ul style="list-style-type: none"> • AD-IDMU - Active Directory Identity Management for UNIX • AD-SFU - Active Directory Services for UNIX • MS-AD-BIS - Active Directory Identity Management for UNIX • RFC-2307 - Schema based on RFC 2307 • Custom schema
servers	array[string]	
session_security	string	<p>Specifies the level of security to be used for LDAP communications:</p> <ul style="list-style-type: none"> • none - no signing or sealing • sign - sign LDAP traffic • seal - seal and sign LDAP traffic
use_start_tls	boolean	<p>Specifies whether or not to use Start TLS over LDAP connections.</p>

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ldap_service

Name	Type	Description
ad_domain	string	This parameter specifies the name of the Active Directory domain used to discover LDAP servers for use by this client. This is mutually exclusive with <code>servers</code> during POST and PATCH.
base_dn	string	Specifies the default base DN for all searches.
base_scope	string	Specifies the default search scope for LDAP queries: <ul style="list-style-type: none"> • base - search the named entry only • onelevel - search all entries immediately below the DN • subtree - search the named DN entry and the entire subtree below the DN
bind_dn	string	Specifies the user that binds to the LDAP servers.
bind_password	string	Specifies the bind password for the LDAP servers.
min_bind_level	string	The minimum bind authentication level. Possible values are: <ul style="list-style-type: none"> • anonymous - anonymous bind • simple - simple bind • sasl - Simple Authentication and Security Layer (SASL) bind
port	integer	The port used to connect to the LDAP Servers.
preferred_ad_servers	array[string]	

Name	Type	Description
schema	string	<p>The name of the schema template used by the SVM.</p> <ul style="list-style-type: none"> • AD-IDMU - Active Directory Identity Management for UNIX • AD-SFU - Active Directory Services for UNIX • MS-AD-BIS - Active Directory Identity Management for UNIX • RFC-2307 - Schema based on RFC 2307 • Custom schema
servers	array[string]	
session_security	string	<p>Specifies the level of security to be used for LDAP communications:</p> <ul style="list-style-type: none"> • none - no signing or sealing • sign - sign LDAP traffic • seal - seal and sign LDAP traffic
svm	svm	
use_start_tls	boolean	Specifies whether or not to use Start TLS over LDAP connections.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage NIS configuration

Security authentication cluster NIS endpoint overview

Overview

NIS servers are used to authenticate user and client computers. NIS domain name and NIS server information is required to configure NIS. This API retrieves and manages NIS server configurations.

Examples

Retrieving cluster NIS information

The cluster NIS GET request retrieves the NIS configuration of the cluster.

The following example shows how a GET request is used to retrieve the cluster NIS configuration:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/nis" -H
"accept: application/hal+json"

# The response:
{
  "domain": "domainA.example.com",
  "servers": [
    "10.10.10.10",
    "example.com"
  ],
  "bound_servers": [
    "10.10.10.10"
  ]
}
```

Creating the cluster NIS configuration

The cluster NIS POST request creates a NIS configuration for the cluster.

The following example shows how a POST request is used to create a cluster NIS configuration:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/nis"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"domain\": \"domainA.example.com\", \"servers\": [
    \"10.10.10.10\", \"example.com\" ]}"
```

Updating the cluster NIS configuration

The cluster NIS PATCH request updates the NIS configuration of the cluster.

The following example shows how to update the domain:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/nis"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"domain\": \"domainC.example.com\", \"servers\": [ \"13.13.13.13\" ]}"
```

The following example shows how to update the server:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/nis"
-H "accept: application/json" -H "Content-Type: application/json" -d "{
  \"servers\": [ \"14.14.14.14\" ]}"
```

Deleting the cluster NIS configuration

The cluster NIS DELETE request deletes the NIS configuration of the cluster.

The following example shows how a DELETE request is used to delete the cluster NIS configuration:

```
# The API:
/security/authentication/cluster/nis

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/authentication/cluster/nis"
-H "accept: application/hal+json"
```

Delete the NIS configuration for the cluster

DELETE /security/authentication/cluster/nis

Introduced In: 9.6

Deletes the NIS configuration of the cluster. NIS can be removed as a source from ns-switch if NIS is not used for lookups.

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the NIS configuration for the cluster

GET /security/authentication/cluster/nis

Introduced In: 9.6

Retrieves the NIS configuration of the cluster. Both NIS domain and servers are displayed by default. The `bound_servers` property indicates the successfully bound NIS servers.

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
bound_servers	array[string]	
domain	string	The NIS domain to which this configuration belongs.
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bound_servers": [
    "string"
  ],
  "domain": "string",
  "servers": [
    "string"
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the NIS configuration for the cluster

PATCH /security/authentication/cluster/nis

Introduced In: 9.6

Both NIS domain and servers can be updated. Domains and servers cannot be empty. Both FQDNs and IP addresses are supported for the 'servers' field. If the domain is updated, NIS servers must also be specified. IPv6 must be enabled if IPv6 family addresses are specified for the `servers` property.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
bound_servers	array[string]	
domain	string	The NIS domain to which this configuration belongs.
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

Example request

```
{
  "bound_servers": [
    "string"
  ],
  "domain": "string",
  "servers": [
    "string"
  ]
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1966253	IPv6 is not enabled in the cluster .
3276964	The NIS domain name or NIS server domain is too long. The maximum supported for domain name is 64 characters and the maximum supported for NIS server domain is 255 characters.
3276933	A maximum of 10 NIS servers can be configured per SVM.
23724109	DNS resolution failed for one or more specified servers.
23724112	DNS resolution failed due to an internal error. Contact technical support if this issue persists.
23724132	DNS resolution failed for all the specified servers.
23724130	Cannot use an IPv6 name server address because there are no IPv6 interfaces

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster_nis_service

Name	Type	Description
bound_servers	array[string]	
domain	string	The NIS domain to which this configuration belongs.
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create the NIS configuration for the cluster

POST /security/authentication/cluster/nis

Introduced In: 9.6

The cluster can have one NIS server configuration. Specify the NIS domain and NIS servers as input. Domain name and servers fields cannot be empty.

Both FQDNs and IP addresses are supported for the `server` property. IPv6 must be enabled if IPv6 family addresses are specified in the `server` property. A maximum of ten NIS servers are supported.

Required properties

- `domain` - NIS domain to which this configuration belongs.
- `servers` - List of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
bound_servers	array[string]	
domain	string	The NIS domain to which this configuration belongs.
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

Example request

```
{
  "bound_servers": [
    "string"
  ],
  "domain": "string",
  "servers": [
    "string"
  ]
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of NIS domain records.
records	array[cluster_nis_service]	

Example response

```
{
  "records": [
    {
      "bound_servers": [
        "string"
      ],
      "domain": "string",
      "servers": [
        "string"
      ]
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1966253	IPv6 is not enabled in the cluster.
3276964	The NIS domain name or NIS server domain is too long. The maximum supported for domain name is 64 characters and the maximum supported for NIS server domain is 255 characters.
3276933	A maximum of 10 NIS servers can be configured per SVM.
13434916	The SVM is in the process of being created. Wait a few minutes, and then try the command again.

Error Code	Description
23724109	DNS resolution failed for one or more specified servers.
23724112	DNS resolution failed due to an internal error. Contact technical support if this issue persists.
23724132	DNS resolution failed for all the specified servers.
23724130	Cannot use an IPv6 name server address because there are no IPv6 interfaces.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster_nis_service

Name	Type	Description
bound_servers	array[string]	
domain	string	The NIS domain to which this configuration belongs.
servers	array[string]	A list of hostnames or IP addresses of NIS servers used by the NIS domain configuration.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage SAML service

Security authentication cluster saml-sp endpoint overview

Overview

This API is used to retrieve and display relevant information pertaining to the SAML service provider configuration in the cluster. The POST request creates a SAML service provider configuration if there is none present. The DELETE request removes the SAML service provider configuration. The PATCH request enables and disables SAML in the cluster. Various responses are shown in the examples below.

+

Examples

Retrieving the SAML service provider configuration in the cluster

The following output shows the SAML service provider configuration in the cluster.

+

```
# The API:
/api/security/authentication/cluster/saml-sp

# The call:
curl -X GET "https://<mgmt-ip>/api/security/authentication/cluster/saml-sp" -H "accept: application/hal+json"

# The response:
{
  "idp_uri": "https://examplelab.customer.com/idp/Metadata",
  "enabled": true,
  "host": "172.21.74.181",
  "certificate": {
    "ca": "cluster1",
    "serial_number": "156F10C3EB4C51C1",
    "common_name": "cluster1"
  },
  "_links": {
    "self": {
      "href": "/api/security/authentication/cluster/saml-sp"
    }
  }
}
```

Creating the SAML service provider configuration

The following output shows how to create a SAML service provider configuration in the cluster.

+

```
# The API:
/api/security/authentication/cluster/saml-sp

# The call:
curl -X POST "https://<mgmt-ip>/api/security/authentication/cluster/saml-sp?return_records=true" -H "accept: application/hal+json" -d '{ "idp_uri": "https://examplelab.customer.com/idp/Metadata", "host": "172.21.74.181", "certificate": { "ca": "cluster1", "serial_number": "156F10C3EB4C51C1" } }'
```

Updating the SAML service provider configuration

The following output shows how to enable a SAML service provider configuration in the cluster.

Disabling the configuration requires the client to be authenticated through SAML prior to performing the operation.

+

```
# The API:
/api/security/authentication/cluster/saml-sp

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/authentication/cluster/saml-sp/" -d '{ "enabled": true }'
```

Deleting the SAML service provider configuration

```
# The API:
/api/security/authentication/cluster/saml-sp

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/authentication/cluster/saml-sp/"
```

Delete a SAML service provider configuration

DELETE /security/authentication/cluster/saml-sp

Introduced In: 9.6

Deletes a SAML service provider configuration.

Response

```
Status: 200, Ok
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
12320803	SAML must be disabled before the configuration can be removed.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a SAML service provider configuration

GET /security/authentication/cluster/saml-sp

Introduced In: 9.6

Retrieves a SAML service provider configuration.

Parameters

Name	Type	In	Required	Description
certificate.serial_number	string	query	False	Filter by certificate.serial_number <ul style="list-style-type: none">Introduced in: 9.7
certificate.ca	string	query	False	Filter by certificate.ca <ul style="list-style-type: none">Introduced in: 9.7

Name	Type	In	Required	Description
certificate.common_name	string	query	False	Filter by certificate.common_name <ul style="list-style-type: none"> Introduced in: 9.7
idp_uri	string	query	False	Filter by idp_uri <ul style="list-style-type: none"> Introduced in: 9.7
host	string	query	False	Filter by host <ul style="list-style-type: none"> Introduced in: 9.7
enabled	boolean	query	False	Filter by enabled <ul style="list-style-type: none"> Introduced in: 9.7
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
certificate	certificate	
enabled	boolean	The SAML service provider is enabled. Valid for PATCH and GET operations only.
host	string	The SAML service provider host.
idp_uri	string	The identity provider (IdP) metadata location. Required for POST operations.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "ca": "string",
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
  },
  "host": "string",
  "idp_uri": "https://idp.example.com/FederationMetadata/2007-06/FederationMetadata.xml"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```


See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

certificate

Name	Type	Description
ca	string	Server certificate issuing certificate authority (CA). This cannot be used with the server certificate common name.
common_name	string	Server certificate common name. This cannot be used with the certificate authority (CA) or serial_number.
serial_number	string	Server certificate serial number. This cannot be used with the server certificate common name.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array [error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Update a SAML service provider configuration

PATCH /security/authentication/cluster/saml-sp

Introduced In: 9.6

Updates a SAML service provider configuration.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
enabled	boolean	The SAML service provider is enabled. Valid for PATCH and GET operations only.

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
12320791	SAML can only be disabled using the console or a SAML-authenticated application.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

certificate

security_saml_sp

Name	Type	Description
enabled	boolean	The SAML service provider is enabled. Valid for PATCH and GET operations only.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a SAML service provider configuration

POST /security/authentication/cluster/saml-sp

Introduced In: 9.6

Creates a SAML service provider configuration. Note that "common_name" is mutually exclusive with

"serial_number" and "ca" in POST. SAML will initially be disabled, requiring a patch to set "enabled" to "true", so that the user has time to complete the setup of the IdP.

Required properties

- idp_uri

Optional properties

- certificate
- enabled
- host

Parameters

Name	Type	In	Required	Description
verify_metadata_server	boolean	query	False	Verify IdP metadata server identity. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
certificate	certificate	
host	string	The SAML service provider host.

Name	Type	Description
idp_uri	string	The identity provider (IdP) metadata location. Required for POST operations.

Example request

```
{
  "certificate": {
    "ca": "string",
    "common_name": "cluster1",
    "serial_number": "1506B24A94F566BA"
  },
  "enabled": null,
  "host": "string",
  "idp_uri": "https://idp.example.com/FederationMetadata/2007-06/FederationMetadata.xml"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

Error Code	Description
12320789	Failed to download data file from specified URI.
12320794	The host parameter provided must be the cluster management interface's IP address. If the cluster management interface is not available, the node management interface's IP address must be used.
12320795	A valid cluster or node management interface IP address must be provided.
12320805	The certificate information provided does not match any installed certificates.
12320806	The certificate information entered does not match any installed certificates.
12320814	An invalid IDP URI has been entered.
12320815	An IDP URI must be an HTTPS or FTPS URI.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

certificate

Name	Type	Description
ca	string	Server certificate issuing certificate authority (CA). This cannot be used with the server certificate common name.
common_name	string	Server certificate common name. This cannot be used with the certificate authority (CA) or serial_number.
serial_number	string	Server certificate serial number. This cannot be used with the server certificate common name.

security_saml_sp

Name	Type	Description
certificate	certificate	
host	string	The SAML service provider host.
idp_uri	string	The identity provider (IdP) metadata location. Required for POST operations.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update user account password

Security authentication password endpoint overview

Overview

This API changes the password for a local user account.

Only cluster administrators with the *"admin"* role can change the password for other cluster or SVM user accounts. If you are not a cluster administrator, you can only change your own password.

Examples

Changing the password of another cluster or SVM user account by a cluster administrator

Specify the user account name and the new password in the body of the POST request. The owner.uuid or owner.name are not required to be specified for a cluster-scoped user account.

For an SVM-scoped account, along with new password and user account name, specify either the SVM name as the owner.name or SVM uuid as the owner.uuid in the body of the POST request. These indicate the SVM for which the user account is created and can be obtained from the response body of a GET request performed on the */api/svm/svms* API.

```
# The API:
POST "/api/security/authentication/password"

# The call to change the password of another cluster user:
curl -X POST "https://<mgmt-ip>/api/security/authentication/password" -d
'{"name":"cluster_user1","password":"hello@1234"}'
```

```
# The call to change the password of another SVM user:
curl -X POST "https://<mgmt-ip>/api/security/authentication/password" -d
'{"owner.name":"svm1","name":"svm_user1","password":"hello@1234"}'
```

Changing the password of an SVM-scoped user



The IP address in the URI must be same as one of the interfaces owned by the SVM.

```
# The API:
POST "/api/security/authentication/password"

# The call:
curl -X POST "https://<SVM-ip>/api/security/authentication/password" -d
'{"name":"svm_user1","password":"new1@1234"}'
```

Update the user account password

POST /security/authentication/password

Introduced In: 9.6

Updates the password for a user account.

Required parameters

- `name` - User account name.
- `password` - New password for the user account.

Optional parameters

- `owner.name` or `owner.uuid` - Name or UUID of the SVM for an SVM-scoped user account.

Related ONTAP commands

- `security login password`

Learn more

- [DOC /security/authentication/password](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
name	string	The user account name whose password is being modified.
owner	owner	Owner name and UUID that uniquely identifies the user account. This field is optional and valid only when a cluster administrator is executing the API to uniquely identify the account whose password is being modified. The "owner" field is not required to be specified for SVM user accounts trying to modify their password.
password	string	The password string

Example request

```
{
  "name": "string",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "password": "string"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
7077918	The password cannot contain the username.
7077919	The minimum length for new password does not meet the policy.
7077920	The new password must have both letters and numbers.
7077921	The minimum number of special characters required do not meet the policy.
7077924	The new password must be different than last N passwords.
7077925	The new password must be different to the old password.
7077940	The password exceeds maximum supported length.
7077941	Defined password composition exceeds the maximum password length of 128 characters.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

owner

Owner name and UUID that uniquely identifies the user account. This field is optional and valid only when a cluster administrator is executing the API to uniquely identify the account whose password is being modified. The "owner" field is not required to be specified for SVM user accounts trying to modify their password.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

account_password

The password object

Name	Type	Description
name	string	The user account name whose password is being modified.
owner	owner	Owner name and UUID that uniquely identifies the user account. This field is optional and valid only when a cluster administrator is executing the API to uniquely identify the account whose password is being modified. The "owner" field is not required to be specified for SVM user accounts trying to modify their password.
password	string	The password string

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage authentication keys (for admins)

Security authentication publickeys endpoint overview

Overview

This API configures the public keys for user accounts.

For secure shell (SSH) access, public-private key pair based authentication is possible by associating the public key with a user account.

Prerequisites:

You must have generated the SSH key.

You must be a cluster or SVM administrator to perform the user's public key.

Examples

Creating a public key for cluster-scoped user accounts

Specify the user account name, public key, index, and comment in the body of the POST request. The owner.uuid or owner.name are not required for a cluster-scoped user account.

```
# The API:
POST "/api/security/authentication/publickey"

# The call
curl -k https://<mgmt-ip>/api/security/authentication/publickeys --request
POST --data '{ "account": "pubuser2", "comment": "Cserver-
Creation", "index": 0, "public_key": "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDWTerk2xXi3Lkrqrm6Q67aNmNDNKgpezwzl9Ap+DxtM
xmocHC52CuwEcJjIUel21uAi0myscaJmC5Y0BzSnYuwcJ3/82kconFUwtK0h6QUTGqbc0vb4MZ
39yIiiItCObWTot3msJYfZB+dgCxCxMZ+0bl9YwZXrWMWP6XDJw/ai/EfWTZHC7e8Xe1mfp+C0c
GMhlxRvgfeGjhZqU85DBfdGD5Tu/67vD13Q+817Jf8iBxgrEFiqirnkWNX5dj+jkM1EVC3s6dY
ciJNBEiMxNkSFGTnkG74S61pvZNFD6mSznV/28hl3zSjwWiuacsP3BelydeG8nEcCRN5Ap7vMd
Vd" }'
```

Creating a public key for SVM-scoped user accounts

For a SVM-scoped account, specify either the SVM name as the owner.name or the SVM UUID as the owner.uuid along with other parameters for the user account. These parameters indicate the SVM that contains the user account for the public key being created and can be obtained from the response body of the GET request performed on the API"/api/svm/svms".

```
# The API:
POST "/api/security/authentication/publickey"

# The call
curl -k https://<mgmt-ip>/api/security/authentication/publickeys --request
POST --data '{ "account": "pubuser4", "comment": "Vserver-
Creation", "index": 0, "owner.uuid": "513a78c7-8c13-11e9-8f78-
005056bbf6ac", "owner.name": "vs0", "public_key": "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDWTerk2xXi3Lkrqrm6Q67aNmNDNKgpezwzl9Ap+DxtM
xmocHC52CuwEcJjIUel21uAi0myscaJmC5Y0BzSnYuwcJ3/82kconFUwtK0h6QUTGqbc0vb4MZ
39yIiiItCObWTot3msJYfZB+dgCxCxMZ+0bl9YwZXrWMWP6XDJw/ai/EfWTZHC7e8Xe1mfp+C0c
GMhlxRvgfeGjhZqU85DBfdGD5Tu/67vD13Q+817Jf8iBxgrEFiqirnkWNX5dj+jkM1EVC3s6dY
ciJNBEiMxNkSFGTnkG74S61pvZNFD6mSznV/28hl3zSjwWiuacsP3BelydeG8nEcCRN5Ap7vMd
Vd" }'
```

Retrieving the configured public key for user accounts

Retrieves all public keys associated with the user accounts or a filtered list (for a specific user account name, a specific SVM and so on) of public keys.

```
# The API:
GET "/api/security/authentication/publickeys"

# The call to retrieve all the user accounts configured in the cluster:
curl -k https://<mgmt-ip>/api/security/authentication/publickeys
```

Retrieve public keys configured for user accounts

GET /security/authentication/publickeys

Introduced In: 9.7

Retrieves the public keys configured for user accounts.

Related ONTAP commands

- `security login publickey show`

Learn more

- [DOC /security/authentication/publickeys](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
scope	string	query	False	Filter by scope
comment	string	query	False	Filter by comment
owner.uuid	string	query	False	Filter by owner.uuid
owner.name	string	query	False	Filter by owner.name
public_key	string	query	False	Filter by public_key
account.name	string	query	False	Filter by account.name
sha_fingerprint	string	query	False	Filter by sha_fingerprint
obfuscated_fingerprint	string	query	False	Filter by obfuscated_fingerprint

Name	Type	In	Required	Description
index	integer	query	False	Filter by index
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
num_records	integer	Number of records
records	array[publickey]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "account": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "joe.smith"
      },
      "comment": "string",
      "obfuscated_fingerprint": "string",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "public_key": "string",
      "scope": "string",
      "sha_fingerprint": "string"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

account_reference

Name	Type	Description
_links	_links	
name	string	User account

owner

Owner name and UUID that uniquely identifies the public key.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

publickey

The public key for the user account (to access SSH).

Name	Type	Description
_links	_links	
account	account_reference	

Name	Type	Description
comment	string	Optional comment for the public key.
index	integer	Index number for the public key (where there are multiple keys for the same account).
obfuscated_fingerprint	string	The obfuscated fingerprint for the public key (READONLY).
owner	owner	Owner name and UUID that uniquely identifies the public key.
public_key	string	The public key
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
sha_fingerprint	string	The SHA fingerprint for the public key (READONLY).

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a public key for a user account

POST /security/authentication/publickeys

Introduced In: 9.7

Creates a public key for a user account.

Required properties

- `owner.uuid` - UUID of the account owner.
- `name` - User account name.
- `index` - Index number for the public key (where there are multiple keys for the same account).
- `public_key` - The publickey details for the creation of the user account.

Related ONTAP commands

- `security login publickey create`

Learn more

- [DOC /security/authentication/publickeys](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
account	account_reference	
comment	string	Optional comment for the public key.
index	integer	Index number for the public key (where there are multiple keys for the same account).
obfuscated_fingerprint	string	The obfuscated fingerprint for the public key (READONLY).

Name	Type	Description
owner	owner	Owner name and UUID that uniquely identifies the public key.
public_key	string	The public key
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
sha_fingerprint	string	The SHA fingerprint for the public key (READONLY).

Example request

```
{
  "account": {
    "name": "joe.smith"
  },
  "comment": "string",
  "obfuscated_fingerprint": "string",
  "owner": {
    "name": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "public_key": "string",
  "scope": "string",
  "sha_fingerprint": "string"
}
```

Response

Status: 201, Created

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

account_reference

Name	Type	Description
name	string	User account

owner

Owner name and UUID that uniquely identifies the public key.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

publickey

The public key for the user account (to access SSH).

Name	Type	Description
account	account_reference	
comment	string	Optional comment for the public key.
index	integer	Index number for the public key (where there are multiple keys for the same account).
obfuscated_fingerprint	string	The obfuscated fingerprint for the public key (READONLY).
owner	owner	Owner name and UUID that uniquely identifies the public key.
public_key	string	The public key

Name	Type	Description
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
sha_fingerprint	string	The SHA fingerprint for the public key (READONLY).

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage authentication keys (end-users)

Security authentication publickeys owner.uuid account.name index endpoint overview

Overview

This API configures the public keys for end-user (non-cluster admin) accounts. Specify the owner UUID, the user account name, and the index in the URI path. The owner UUID corresponds to the UUID of the SVM containing the user account associated with the public key and can be obtained from the response body of the GET request performed on the API `/api/svm/svms`.

The index value corresponds to the public key that needs to be modified or deleted (it is possible to create more than one public key for the same user account).

Examples

Retrieving the specific configured public key for user accounts

```
# The API:
GET
"/api/security/authentication/publickeys/{owner.uuid}/{account.name}/{index}"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/publickeys/513a78c7-8c13-11e9-8f78-005056bbf6ac/pubuser4/0
```

Updating the public key and comment for user accounts

```
# The API:
PATCH
"/api/security/authentication/publickeys/{owner.uuid}/{account.name}/{index}"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/publickeys/d49de271-8c11-11e9-8f78-005056bbf6ac/pubuser1/0 --request PATCH --data '{
"comment": "Cserver-modification", "public_key": "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQCMsLP/FeiT1J4Fb4GNVO4ioa1NIUHWeG08+anDbFke3J
cFT5JqBn0QZiG0uF0bqepken/moVKZg8iQng1arjP4ULhhje/LwDuUbaB7kvtPL2gyzAX1qFYn
BJ5R1LXja25Z4xeeaXUBJjhUmvpfque0TxbvpaG5V9rFTzVg9ccjBnkBchg3EkhF4VtHmrZNpT
DAUOBaz69FRYXYZ2ExoCHWqElHBJep9D0DLN0XtzQA0IF9hJck6xja5RcAQ6f9pLMCol9vJiqp
cBAjkUmg1qH5ZNHsgDQ7dtGNGJw45zqXHPAy9z8yKJuIsdK2/4iVYLDL8mlHFEIgeADn6OSxui
j1" }'
```

Deleting the public key for user accounts

```
# The API:
DELETE
"/api/security/authentication/publickeys/{owner.uuid}/{account.name}/{index}"

# The call:
curl -k https://<mgmt-ip>/api/security/authentication/publickeys/d49de271-8c11-11e9-8f78-005056bbf6ac/pubuser1/0 --request DELETE
```

Delete a public key for a user account

```
DELETE /security/authentication/publickeys/{owner.uuid}/{account.name}/{index}
```

Introduced In: 9.7

Deletes the public key for a user account.

Related ONTAP commands

- `security login publickey delete`

Learn more

- [DOC /security/authentication/publickeys/{owner.uuid}/{account.name}/{index}](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Account owner UUID
account.name	string	path	True	User account name
index	integer	path	True	Index number for the public key (where there are multiple keys for the same account). <ul style="list-style-type: none">• Max value: 99• Min value: 0

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve public keys configured for a user account

GET /security/authentication/publickeys/{owner.uuid}/{account.name}/{index}

Introduced In: 9.7

Retrieves the public keys configured for a user account.

Related ONTAP commands

- `security login publickey show`

Learn more

- [DOC /security/authentication/publickeys/{owner.uuid}/{account.name}/{index}](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Account owner UUID
account.name	string	path	True	User account name
index	integer	path	True	Index number for the public key (where there are multiple keys for the same account). <ul style="list-style-type: none">• Max value: 99• Min value: 0
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Max value: 120 • Min value: 0 • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
account	account_reference	
comment	string	Optional comment for the public key.
index	integer	Index number for the public key (where there are multiple keys for the same account).
obfuscated_fingerprint	string	The obfuscated fingerprint for the public key (READONLY).
owner	owner	Owner name and UUID that uniquely identifies the public key.
public_key	string	The public key

Name	Type	Description
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
sha_fingerprint	string	The SHA fingerprint for the public key (READONLY).

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "account": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "joe.smith"
  },
  "comment": "string",
  "obfuscated_fingerprint": "string",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "public_key": "string",
  "scope": "string",
  "sha_fingerprint": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

account_reference

Name	Type	Description
_links	_links	
name	string	User account

owner

Owner name and UUID that uniquely identifies the public key.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Update a public key for a user account

PATCH /security/authentication/publickeys/{owner.uuid}/{account.name}/{index}

Introduced In: 9.7

Updates the public key for a user account.

Related ONTAP commands

- `security login publickey modify`

Learn more

- [DOC /security/authentication/publickeys/{owner.uuid}/{account.name}/{index}](#)
- [DOC /security/accounts](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Account owner UUID
account.name	string	path	True	User account name
index	integer	path	True	Index number for the public key (where there are multiple keys for the same account). <ul style="list-style-type: none"> • Max value: 99 • Min value: 0

Request Body

Name	Type	Description
comment	string	Optional comment for the public key.

Name	Type	Description
obfuscated_fingerprint	string	The obfuscated fingerprint for the public key (READONLY).
public_key	string	The public key
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
sha_fingerprint	string	The SHA fingerprint for the public key (READONLY).

Example request

```
{
  "comment": "string",
  "obfuscated_fingerprint": "string",
  "public_key": "string",
  "scope": "string",
  "sha_fingerprint": "string"
}
```

Response

```
Status: 200, Ok
```

Error

```
Status: Default, Error
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

account_reference

Name	Type	Description
name	string	User account

owner

Owner name and UUID that uniquely identifies the public key.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

publickey

The public key for the user account (to access SSH).

Name	Type	Description
comment	string	Optional comment for the public key.
obfuscated_fingerprint	string	The obfuscated fingerprint for the public key (READONLY).
public_key	string	The public key
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.
sha_fingerprint	string	The SHA fingerprint for the public key (READONLY).

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage Azure Key Vaults

Security azure-key-vaults endpoint overview

Overview

Azure Key Vault (AKV) is a cloud key management service (KMS) that provides a secure store for secrets. This feature

allows ONTAP to securely store its encryption keys using AKV.

In order to use AKV with ONTAP, you must first deploy an Azure application with the appropriate access to an AKV and then provide

ONTAP with the necessary details, such as key vault name, application ID so that ONTAP can communicate with the deployed Azure application.

Examples

Creating an AKV for a cluster

The example AKV is configured at the cluster-scope. Note the *return_records=true* query parameter is used to obtain the newly created key manager configuration.

```
# The API:
POST /api/security/azure-key-vaults

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/azure-key-
vaults?return_records=true' -H 'accept: application/hal+json' -d "{
  \"client_id\": \"client1\", \"tenant_id\": \"tenant1\", \"name\":
  \"https://mykeyvault.azure.vault.net/\", \"key_id\": \"https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74\",
  \"client_secret\" : \"myclientPwd\" }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
      "client_id": "client1",
      "tenant_id": "tenant1",
      "name": "https://mykeyvault.azure.vault.net/",
      "key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
      "_links": {
        "self": {
          "href": "/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-
005056bbeba5"
        }
      }
    }
  ]
}
```

Creating an AKV for an SVM

The example AKV is configured for a specific SVM. Note the *return_records=true* query parameter is used to obtain the newly created key manager configuration.

```
# The API:
POST /api/security/azure-key-vaults

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/azure-key-
vaults?return_records=true' -H 'accept: application/hal+json' -d "{
  \"svm\": { \"uuid\": \"4f7abf4c-9a07-11ea-8d52-005056bbeba5\" },
  \"client_id\": \"client1\", \"tenant_id\": \"tenant1\", \"name\":
  \"https://mykeyvault.azure.vault.net/\", \"key_id\": \"https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74\",
  \"client_secret\" : \"myclientPwd\" }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "024cd3cf-9a08-11ea-8d52-005056bbeba5",
      "svm": {
        "uuid": "4f7abf4c-9a07-11ea-8d52-005056bbeba5",
        "name": "vs0"
      },
      "client_id": "client1",
      "tenant_id": "tenant1",
      "name": "https://mykeyvault.azure.vault.net/",
      "key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
      "_links": {
        "self": {
          "href": "/api/security/azure-key-vaults/024cd3cf-9a08-11ea-8d52-
005056bbeba5"
        }
      }
    }
  ]
}
```

Retrieving the AKVs configured for all clusters and SVMs

The following example shows how to retrieve all configured AKVs along with their configurations.

```
# The API:
GET /api/security/azure-key-vaults
```



```
# The call:
curl -X GET 'https://<mgmt-ip>/api/security/azure-key-vaults?fields=*'

# The response:
{
  "records": [
    {
      "uuid": "024cd3cf-9a08-11ea-8d52-005056bbeba5",
      "scope": "svm",
      "svm": {
        "uuid": "4f7abf4c-9a07-11ea-8d52-005056bbeba5",
        "name": "vs0"
      },
      "client_id": "client1",
      "tenant_id": "tenant1",
      "name": "https://mykeyvault.azure.vault.net/",
      "key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
      "state": {
        "cluster_state": true,
        "message": "",
        "code": 0
      },
      "_links": {
        "self": {
          "href": "/api/security/azure-key-vaults/024cd3cf-9a08-11ea-8d52-
005056bbeba5"
        }
      }
    },
    {
      "uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
      "scope": "cluster",
      "client_id": "client1",
      "tenant_id": "tenant1",
      "name": "https://mykeyvault.azure.vault.net/",
      "key_id": "https://keyvault-
test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
      "state": {
        "cluster_state": true,
        "message": "",
        "code": 0
      },
      "_links": {
        "self": {
```

```
      "href": "/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-005056bbeba5"
    }
  }
},
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/security/azure-key-vaults?fields=*"
  }
}
}
```

Retrieving the AKV configured for a specific SVM

The following example retrieves a configured AKV for a specific SVM.

```
# The API:
GET /api/security/azure-key-vaults

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-005056bbeba5?fields=*'

# The response:
{
  "uuid": "85619643-9a06-11ea-8d52-005056bbeba5",
  "scope": "cluster",
  "client_id": "client1",
  "tenant_id": "tenant1",
  "name": "https://mykeyvault.azure.vault.net/",
  "key_id": "https://keyvault-test.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
  "state": {
    "cluster_state": true,
    "message": "",
    "code": 0
  },
  "_links": {
    "self": {
      "href": "/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-005056bbeba5"
    }
  }
}
```

Updating the client password of a specific SVM

The following example updates the client password of a configured AKV for a specific SVM.

```
# The API:
PATCH /api/security/azure-key-vaults

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-005056bbeba5' -d '{"client_secret": "newSecret"}'
```

Deleting an AKV configuration for a specific SVM

The following example deletes a configured AKV for a specific SVM.

```
# The API:
DELETE /api/security/azure-key-vaults

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-005056bbeba5'
```

Restoring the keys for a specific SVM configured with an AKV

The following example restores all the keys of a specific SVM configured with an AKV.

```
# The API:
POST security/azure-key-vaults/{azure_key_vault.uuid}/restore

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/azure-key-vaults/85619643-9a06-11ea-8d52-005056bbeba5/restore'

# The response:
{
  "job": {
    "uuid": "6ab6946f-9a0c-11ea-8d52-005056bbeba5",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6ab6946f-9a0c-11ea-8d52-005056bbeba5"
      }
    }
  }
}
```

Retrieve AKVs configured for all clusters and SVMs

GET /security/azure-key-vaults

Introduced In: 9.8

Retrieves AKVs configured for all clusters and SVMs.

Related ONTAP commands

- `security key-manager external azure show`

Parameters

Name	Type	In	Required	Description
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
tenant_id	string	query	False	Filter by tenant_id
state.code	integer	query	False	Filter by state.code
state.message	string	query	False	Filter by state.message
state.cluster_state	boolean	query	False	Filter by state.cluster_state
scope	string	query	False	Filter by scope
name	string	query	False	Filter by name
key_id	string	query	False	Filter by key_id
uuid	string	query	False	Filter by uuid
client_id	string	query	False	Filter by client_id
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[azure_key_vault]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
      "key_id":
"https://keyvault1.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c5954
0e2a74",
      "name": "https://kmip-akv-keyvault.vault.azure.net/",
      "scope": "string",
      "state": {
        "code": "346758",
        "message": "AKV key protection is unavailable in following
nodes - node1, node2."
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzzz",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

state

Indicates whether the AKV key protection is available cluster wide.

Name	Type	Description
cluster_state	boolean	Set to true when AKV key protection is available on all nodes of the cluster.
code	integer	Code corresponding to the status message. Returns a 0 if AKV key protection is available in all nodes of the cluster.
message	string	Error message set when cluster availability is false.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

azure_key_vault

Name	Type	Description
_links	_links	
client_id	string	Application client ID of the deployed Azure application with appropriate access to an AKV.
key_id	string	Key Identifier of AKV key encryption key.
name	string	<p>Name of the deployed AKV that will be used by ONTAP for storing keys.</p> <ul style="list-style-type: none"> • example: https://kmip-akv-keyvault.vault.azure.net/ • format: uri • Introduced in: 9.8
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
state	state	Indicates whether the AKV key protection is available cluster wide.
svm	svm	
tenant_id	string	Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV.
uuid	string	A unique identifier for the Azure Key Vault (AKV).

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create an AKV configuration for all clusters and SVMs

POST /security/azure-key-vaults

Introduced In: 9.8

Configures the AKV configuration for all clusters and SVMs.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create a AKV.
- `client_id` - Application (client) ID of the deployed Azure application with appropriate access to an AKV.
- `tenant_id` - Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV.
- `client_secret` - Password used by the application to prove its identity to AKV.
- `key_id` - Key Identifier of AKV encryption key.
- `name` - Name of the deployed AKV used by ONTAP for storing keys.

Related ONTAP commands

- `security key-manager external azure enable`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
client_id	string	Application client ID of the deployed Azure application with appropriate access to an AKV.
client_secret	string	Password used by the application to prove its identity to AKV.
key_id	string	Key Identifier of AKV key encryption key.
name	string	<p>Name of the deployed AKV that will be used by ONTAP for storing keys.</p> <ul style="list-style-type: none"> • example: https://kmip-akv-keyvault.vault.azure.net/ • format: uri • Introduced in: 9.8
svm	svm	
tenant_id	string	Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV.
uuid	string	A unique identifier for the Azure Key Vault (AKV).

Example request

```
{
  "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
  "client_secret": "abcdef",
  "key_id":
  "https://keyvault1.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c5954
  0e2a74",
  "name": "https://kmip-akv-keyvault.vault.azure.net/",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[azure_key_vault]	

Example response

```
{
  "records": [
    {
      "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
      "client_secret": "abcdef",
      "key_id":
        "https://keyvault1.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c5954
        0e2a74",
      "name": "https://kmip-akv-keyvault.vault.azure.net/",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
3735553	Failed to create self-signed certificate.
3735664	The specified key size is not supported in FIPS mode.
3735665	The specified hash function is not supported in FIPS mode.

Error Code	Description
3735700	The specified key size is not supported.
52559972	The certificates start date is later than the current date.
65537500	A key manager has already been configured for this SVM.
65537503	Passwords do not match.
65537504	Internal error. Failed to store configuration in internal database.
65537505	One or more volume encryption keys of the given SVM are stored on a key manager configured for the admin SVM.
65537506	AKV is not supported in MetroCluster configurations.
65537512	AKV cannot be configured for the given SVM as not all nodes in the cluster have the azure.key.vault capability needed to enable the feature.
65537514	Failed to check or update the azure.key.vault feature capability.
65537518	Failed to find an interface with Cluster role.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

state

Indicates whether the AKV key protection is available cluster wide.

Name	Type	Description
cluster_state	boolean	Set to true when AKV key protection is available on all nodes of the cluster.
code	integer	Code corresponding to the status message. Returns a 0 if AKV key protection is available in all nodes of the cluster.
message	string	Error message set when cluster availability is false.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

azure_key_vault

Name	Type	Description
client_id	string	Application client ID of the deployed Azure application with appropriate access to an AKV.
client_secret	string	Password used by the application to prove its identity to AKV.
key_id	string	Key Identifier of AKV key encryption key.

Name	Type	Description
name	string	Name of the deployed AKV that will be used by ONTAP for storing keys. <ul style="list-style-type: none"> • example: https://kmip-akv-keyvault.vault.azure.net/ • format: uri • Introduced in: 9.8
svm	svm	
tenant_id	string	Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV.
uuid	string	A unique identifier for the Azure Key Vault (AKV).

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Restore keys for an SVM

POST /security/azure-key-vaults/{azure_key_vault.uuid}/restore

Introduced In: 9.8

Restore the keys for an SVM from a configured AKV.

Related ONTAP commands

- `security key-manager external azure restore`

Parameters

Name	Type	In	Required	Description
azure_key_vault.uuid	string	path	True	UUID of the existing AKV configuration.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none">• Default value:

Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
65537511	AKV is not configured for the given SVM.
65537515	Failed to restore keys on some nodes in the cluster.

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete an AKV configuration

DELETE /security/azure-key-vaults/{uuid}

Introduced In: 9.8

Deletes an AKV configuration.

Related ONTAP commands

- `security key-manager external azure disable`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	AKV UUID

Response

```
Status: 200, Ok
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
65536242	One or more Storage Encryption devices are assigned an authentication key.
65536817	Internal error. Failed to determine if key manager is safe to disable.
65536827	Internal error. Failed to determine if the given SVM has any encrypted volumes.
65536834	Internal error. Failed to get existing key-server details for the given SVM.
65536867	Volume encryption keys (VEK) for one or more encrypted volumes are stored on the key manager configured for the given SVM.
65536883	Internal error. Volume encryption key is missing for a volume.
65536884	Internal error. Volume encryption key is invalid for a volume.
65536924	Cannot remove key manager that still contains one or more NSE authentication keys.
65537511	AKV is not configured for the given SVM.
196608080	One or more nodes in the cluster have the root volume encrypted using NVE (NetApp Volume Encryption).

Error Code	Description
196608301	Internal error. Failed to get encryption type.
196608305	NAE aggregates found in the cluster.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the AKV configuration for an SVM specified by the UUID

GET /security/azure-key-vaults/{uuid}

Introduced In: 9.8

Retrieves the AKV configuration for the SVM specified by the UUID.

Related ONTAP commands

- security key-manager external azure show

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	AKV UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
client_id	string	Application client ID of the deployed Azure application with appropriate access to an AKV.
key_id	string	Key Identifier of AKV key encryption key.
name	string	<p>Name of the deployed AKV that will be used by ONTAP for storing keys.</p> <ul style="list-style-type: none"> • example: https://kmip-akv-keyvault.vault.azure.net/ • format: uri • Introduced in: 9.8
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
state	state	Indicates whether the AKV key protection is available cluster wide.
svm	svm	
tenant_id	string	Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV.
uuid	string	A unique identifier for the Azure Key Vault (AKV).

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaa",
  "key_id":
  "https://keyvault1.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
  "name": "https://kmip-akv-keyvault.vault.azure.net/",
  "scope": "string",
  "state": {
    "code": "346758",
    "message": "AKV key protection is unavailable in following nodes - node1, node2."
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

state

Indicates whether the AKV key protection is available cluster wide.

Name	Type	Description
cluster_state	boolean	Set to true when AKV key protection is available on all nodes of the cluster.
code	integer	Code corresponding to the status message. Returns a 0 if AKV key protection is available in all nodes of the cluster.
message	string	Error message set when cluster availability is false.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the AKV configuration

PATCH /security/azure-key-vaults/{uuid}

Introduced In: 9.8

Updates the AKV configuration.

Required properties

- `client_secret` - New password used to prove the application's identity to the AKV.

Related ONTAP commands

- `security key-manager external azure modify`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	AKV UUID

Request Body

Name	Type	Description
client_id	string	Application client ID of the deployed Azure application with appropriate access to an AKV.
client_secret	string	Password used by the application to prove its identity to AKV.
key_id	string	Key Identifier of AKV key encryption key.

Name	Type	Description
name	string	Name of the deployed AKV that will be used by ONTAP for storing keys. <ul style="list-style-type: none"> • example: https://kmip-akv-keyvault.vault.azure.net/ • format: uri • Introduced in: 9.8
svm	svm	
tenant_id	string	Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV.
uuid	string	A unique identifier for the Azure Key Vault (AKV).

Example request

```
{
  "client_id": "aaaaaaaa-bbbb-aaaa-bbbb-aaaaaaaaaaaaa",
  "client_secret": "abcdef",
  "key_id":
    "https://keyvault1.vault.azure.net/keys/key1/a8e619fd8f234db3b0b95c59540e2a74",
  "name": "https://kmip-akv-keyvault.vault.azure.net/",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tenant_id": "zzzzzzzz-yyyy-zzzz-yyyy-zzzzzzzzzzzz",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
65537503	Passwords do not match.
65537504	Internal error. Failed to store configuration in internal database.
65537511	AKV is not configured for the given SVM.
65537517	The field "new_client_secret" must be specified.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

state

Indicates whether the AKV key protection is available cluster wide.

Name	Type	Description
cluster_state	boolean	Set to true when AKV key protection is available on all nodes of the cluster.
code	integer	Code corresponding to the status message. Returns a 0 if AKV key protection is available in all nodes of the cluster.
message	string	Error message set when cluster availability is false.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

azure_key_vault

Name	Type	Description
client_id	string	Application client ID of the deployed Azure application with appropriate access to an AKV.
client_secret	string	Password used by the application to prove its identity to AKV.
key_id	string	Key Identifier of AKV key encryption key.

Name	Type	Description
name	string	Name of the deployed AKV that will be used by ONTAP for storing keys. <ul style="list-style-type: none"> • example: https://kmip-akv-keyvault.vault.azure.net/ • format: uri • Introduced in: 9.8
svm	svm	
tenant_id	string	Directory (tenant) ID of the deployed Azure application with appropriate access to an AKV.
uuid	string	A unique identifier for the Azure Key Vault (AKV).

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage security certificates

Security certificates endpoint overview

Overview

This API displays security certificate information and manages the certificates in ONTAP.

Installing certificates in ONTAP

The security certificates GET request retrieves all of the certificates in the cluster.

Examples

Retrieving all certificates installed in the cluster with their common-names

```
# The API:
/api/security/certificates

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/certificates?fields=common_name" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "svm": {
        "name": "vs0"
      },
      "uuid": "dad2363b-8ac0-11e8-9058-005056b482fc",
      "common_name": "vs0",
      "_links": {
        "self": {
          "href": "/api/security/certificates/dad2363b-8ac0-11e8-9058-
005056b482fc"
        }
      }
    },
    {
      "uuid": "1941e048-8ac1-11e8-9058-005056b482fc",
      "common_name": "ROOT",
      "_links": {
        "self": {
          "href": "/api/security/certificates/1941e048-8ac1-11e8-9058-
005056b482fc"
        }
      }
    },
    {
      "uuid": "5a3a77a8-892d-11e8-b7da-005056b482fc",
```

```
    "common_name": "gshancluster-4",
    "_links": {
      "self": {
        "href": "/api/security/certificates/5a3a77a8-892d-11e8-b7da-005056b482fc"
      }
    }
  },
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/security/certificates?fields=common_name"
    }
  }
}
```

Retrieving all certificates installed at cluster-scope with their common-names

```
# The API:
/api/security/certificates

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/certificates?scope=cluster&fields=common_name" -H
"accept: application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "1941e048-8ac1-11e8-9058-005056b482fc",
      "scope": "cluster",
      "common_name": "ROOT",
      "_links": {
        "self": {
          "href": "/api/security/certificates/1941e048-8ac1-11e8-9058-
005056b482fc"
        }
      }
    },
    {
      "uuid": "5a3a77a8-892d-11e8-b7da-005056b482fc",
      "scope": "cluster",
      "common_name": "gshancluster-4",
      "_links": {
        "self": {
          "href": "/api/security/certificates/5a3a77a8-892d-11e8-b7da-
005056b482fc"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/security/certificates?scope=cluster&fields=common_name"
    }
  }
}
```

Retrieving all certificates installed on a specific SVM with their common-names

```
# The API:
/api/security/certificates

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/certificates?svm.name=vs0&fields=common_name" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "svm": {
        "name": "vs0"
      },
      "uuid": "dad2363b-8ac0-11e8-9058-005056b482fc",
      "common_name": "vs0",
      "_links": {
        "self": {
          "href": "/api/security/certificates/dad2363b-8ac0-11e8-9058-
005056b482fc"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/certificates?svm.name=vs0&fields=common_name"
    }
  }
}
```

Retrieving a certificate using its UUID for all fields

```
# The API:
/api/security/certificates/{uuid}
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/security/certificates/dad2363b-8ac0-11e8-9058-005056b482fc?fields=*" -H "accept: application/hal+json"

# The response:
{
  "svm": {
    "uuid": "dad2363b-8ac0-11e8-9058-005056b482fc",
    "name": "vs0"
  },
  "uuid": "dad2363b-8ac0-11e8-9058-005056b482fc",
  "scope": "svm",
  "type": "server",
  "common_name": "vs0",
  "serial_number": "15428D45CF81CF56",
  "ca": "vs0",
  "hash_function": "sha256",
  "key_size": 2048,
  "expiry_time": "2019-07-18T15:29:14-04:00",
  "public_certificate": "-----BEGIN CERTIFICATE-----
\nMIIDQjCCAiqgAwIBAgIIFUKNRc+Bz1YwDQYJKoZIhvcNAQELBQAwwGZEMMAoGA1UE\nnAxMDdnMwMQswCQYDVQQGEwJVUzAeFw0xODA3MTgxOTI1MTRaFw0xOTA3MTgxOTI1\nnMTRaMBsxDDAKBgNVBAMTA3ZzMDELMAkGA1UEBhMCVVMwggeEiMA0GCSqGSIsb3DQE\nB\nnAQUA4IBDWAawggEKAAIBAQCqFQB27th2ACOMJvWgLhlXRzobSb2ZTQfO561faXQ3\n\nnIbiT+rnrWXetd/s2+iCv9ld9LW0NOMP3MN2f3SFbyze3dl7WrnVbjLmYuI9MfOxs\n\nnfma+Bh6gpap5Yn2YddqoV6rfNGAuUveNLArNl8wODk/mpawPEQ93QSa1zfglgnOH\n\nnRFRyqiSYT06X5g6RbUuEl4LTGXspz+plU46Za0i6QyxvtvZ4bneibffFXN3IigpqI6\n\nnTGUV8R/J3Ps338VxVmSO9ZXBZmvbcJVoySYNICl/oif3fgPZlnBv0tbSwqg4FoZO/\n\nnWT+XHGHLeP6cr/Aqg7u6C4RfqBCwzB/XFKDIqnMAQKDBAGMBAAGjgYkwgYYwDAYD\n\nnVR0TBAUwAwEB/zALBgNVHQ8EBAMCAQYWwHQYDVR0OBByEFN/AnH8qLxocTtumNHIn\n\nnEN4IFIDBMEoGA1UdIwRDMEGAFN/AnH8qLxocTtumNHInEN4IFIDBoR+kHTAbMQww\n\nnCgYDVQQDEwN2czAxZCZAJBgNVBAYTA1VTgggVQo1Fz4HPVjANBgkqhkiG9w0BAQSF\n\nnAAOCAQEAA0pUEepdeQnd2Amwg8UFyxayb8eu3E6dlptvtyp+xtjhIC7Dh95CVXhy\n\n\nnkJS3TSu60PGR/b2vc3MZtAUpcL4ceD8XntKPQgBlqoB4bRogCelTnlGswRXDX5TS\n\n\nngMVrRjaWTBF7ikt4UjR05rSxcDGplQRqjnOthqi+yPT+29+8a4Uu6J+3Kdrflj4p\n\n\nnlnSWpuB9EyxtuCILNQXA2ncH7YKtoeNtChKCchhvPcoTy6Opma6UQN5UMxstkvGT\n\n\nnVGaN5TlRWv0yiqPXIQblSqXi/uQsuRPCHdu7+KWRFn08USa6QVo2mDs9P7R9dd0K\n\n\nn9QAsTjTOF9PlAKgNxGoOJl2y0+48AA==\n\n\n-----END CERTIFICATE-----
\n",
  "_links": {
    "self": {
      "href": "/api/security/certificates/dad2363b-8ac0-11e8-9058-005056b482fc"
    }
  }
}
```

Creating a certificate in a cluster

These certificates can be used to help administrators enable certificate-based authentication and to enable SSL-based communication to the cluster.

```
# The API:
/api/security/certificates

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
  \"common_name\": \"TEST-SERVER\",  \"type\": \"server\"  }"
```

Installing a certificate in a cluster

These certificates can be used to help administrators enable certificate-based authentication and to enable-SSL based communication to the cluster.

```
# The API:
/api/security/certificates

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"type\":
\"server_ca\", \"public_certificate\": \"-----BEGIN CERTIFICATE-----
\\nMIIIFYDCCA0igAwIBAgIQCgFCgAAAAUjyESlAAAAAaJANBgkqhkiG9w0BAQsFADBKMqswCQYD
VQQG\\nEwJVUzESMBAGA1UEChMJSWRlbnRydXN0MScwJQYDVQQDEx5JZGVuVHJlc3QgQ29tbWVy
Y2lhbCBS\\nb290IENBIDEwHhcNMTQwMTE2MTgxMjIzWhcNMzQwMTE2MTgxMjIzWjBKMqswCQYD
VQQGEwJVUzES\\nMBAGA1UEChMJSWRlbnRydXN0MScwJQYDVQQDEx5JZGVuVHJlc3QgQ29tbWVy
Y2lhbCBSb290IENB\\nIDEwggIiMA0GCSqGSIb3DQEBAQUAA4ICDwAwggIKAoICAQCnUBneP5k9
1DNG8W9RYYKYqU+PZ4ld\\nhNlT3Qwo2dfw/66VQ3KZ+bVdfIrBQuExUHRgQl8zZshq0PirKle
hm7zCYofWjK9ouuU+ehcCuz/\\nmNKvcb00U59Oh++SvL3sTzIwiEsXXlfEU8L2ApeN2WIrvyQf
Yo3fw7gpS0l4PJNgiCL8mdo2yMKi\\nlCxUAGclbnO/AljwpN3lsKImesrgNqUZFvX9t++uP0D1
bVoE/c40yiTcdCMbXTMTEl3EASX2MN0C\\nXZ/g1Ue9tOsobobTJSdifWwLziuQkkORiT0/Br4sO
dBeo0XKIanoBScy0RnnGF7HamB4HWfp1IYVl\\n3ZBWzvurpWCdxJ35UrCLvYf5jysjCiN2O/cz
4ckA82n5S6LgTrx+kzmEB/dEcH7+B1rlsazRGMzy\\nNeVJSQjKVsk9+w8YfYs7wRPCTY/JTw43
6R+hDmrfYi7LNQZReSzIJTj0+kuniVyc0uMNOYzKdHzV\\nWYfCP04MXFL0PfdSgvHqo6z9STQa
KPNBiDoT7uje/5kdX7rL6B7yuVBgwDHTc+XvvqDtMwt0viAg\\nxGds8AgDelWAF0ZOlqf0Hj7h
9tgJ4TNkK2PXMl6f+cb7D3hvl7yTmvmcEpB4eoCHFddydJxVdHix\\nuuFucAS6T6C6aMN7/zHw
cz09lCqxCOEOoP5NiGVreTO0lwIDAQABo0IwQDAOBgNVHQ8BAf8EBAMC\\nAQYwDwYDVR0TAQH/
BAUwAwEB/zAdBgNVHQ4EFgQU7UQZwNPwBovupHu+QucmVMiONnYwDQYJKoZI\\nhvcNAQELBQAD
ggIBAA2ukDL2pkt8RHYZYR4nKMleVO8lvOMIkPkp165oCOGUAFjvLi5+U1KMtlwH\\n6oi6mYtQ
lNeCgN9hCQCTrQ0U5s7B8jeUeLBfnLOic7iPBZM4zy0+sLj7wM+x8uwtLRvM7Kqas6pg\\nghst
O8OEPVeKlh6cdbjTMM1gCIOQ045U8UlmwF10A0Cj7oV+wh93nAbowacYXVKV7cndJZ5t+qnt\\n
ozo00F172ulQ8zW/7esUTTHHYPTa8Yec4kjixsU3+wYQ+nVZZjFHKdp2mhZpgq7vmrlR94gjmm
mV\\nYjzlVYA211QC//G5Xc7UI2/YRYRKW2XviQzdFKcgyxilJbQN+QHwotL0AMh0jqEqSI5l2x
PE4iUX\\nfeu+hlsXIFRRk0pTAwvsXcoz7WL9RccvW9xYoIA55vrX/hMUpu09lEpCdNTDd1lzzY
9GvlU47/ro\\nkTLq1lgEIt44w8y8bckzOmoKaT+gyOpyj4xjhi09bTyWnpXgSUyqorkqG5w2gX
jtw+hG4iZZRHUe\\n2XWJUc0QhJ1hYMTd+ZciTY6Y5uN/9lu7rs3KSoFrXgvzUeF0K+l+J6fZmU
lO+KWA2yUPHGNiiskz\\nZ2s8EIPGrd6ozRaOjfAHN3Gf8qv8QfXBi+wAN10J5U6A7/qxXDgGpR
tK4dw4LTzcqx+QGtVKno7R\\ncGzM7vRX+Bi6hG6H\\n-----END CERTIFICATE-----\\n\"
}"
```

Installing a certificate on a specific SVM

```
# The API:
/api/security/certificates

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates" -H "accept:
application/json" -H "Content-Type: application/json" -d "{  \"svm\" : {
  \"name\" : \"vs0\" }, \"type\": \"server_ca\", \"public_certificate\":
  \"-----BEGIN CERTIFICATE-----
\nMIIIFYDCCA0igAwIBAgIQCgFCgAAAAUjyES1AAAAjANBgkqhkiG9w0BAQsFADBKMQswCQYD
VQQG\nEwJVUzESMBAGA1UEChMJSWRlb1RydXN0MScwJQYDVQQDEx5JZGVuVHJ1c3QgQ29tbWVy
Y2lhbCBS\nb290IENBIDEwHhcNMTQwMTE2MTgxmjIzWhcNMzQwMTE2MTgxmjIzWjBKMQswCQYD
VQQGEwJVUzES\nMBAGA1UEChMJSWRlb1RydXN0MScwJQYDVQQDEx5JZGVuVHJ1c3QgQ29tbWVy
Y2lhbCBSb290IENB\nIDewggIiMA0GCSqGSIb3DQEBAQUAA4ICDwAwggIKAoICAQCnUBneP5k9
1DNG8W9RYYKyqU+PZ4ld\nhNlT3Qwo2dfw/66VQ3KZ+bVdfIrBQuExUHTRgQ18zZshq0PirK1e
hm7zCYofWjK9ouuU+ehcCuz/\nmNKvcb00U59Oh++SvL3sTzIwiEsXXlFEU8L2ApeN2WIrVYQf
Yo3fw7gps014PJNgiCL8mdo2yMKi\nlCcxUAGclbnO/AljwpN3lsKImesrgNqUZFvX9t++uP0D1
bVoE/c40yiTcdCMbXTMTEl3EASX2MN0C\nXZ/g1Ue9tOsbobtJSdifWwLziuQkkORiT0/Br4sO
dBeo0XKlIanoBScy0RnnGF7HamB4HWfp1IYVl\n3ZBWzvurpWCdxJ35UrClvYf5jysjCiN2O/cz
4ckA82n5S6LgTrx+kzmEB/dEcH7+B1rlsazRGMzy\nNeVJSQjKVsk9+w8YfYs7wRPCTY/JTw43
6R+hDmrfYi7LNQZReSzIJTj0+kuniVyc0uMNOYZKdHzV\nWYfCP04MXFL0PfdSgvHqo6z9STQa
KPNBiDoT7uje/5kdX7rL6B7yuVBgwDHTc+XvvqDtMwt0viAg\nxGds8AgDelWaf0ZOlqf0Hj7h
9tgJ4TNkK2PXMl6f+cB7D3hvl7yTmvmcEpB4eoCHFddyJxVdHix\nnuuFucAS6T6C6aMN7/zHw
cz09lCqxC0EOoP5NiGVreTO0lwIDAQABo0IwQDAOBgNVHQ8BAf8EBAMC\naQYwDwYDVR0TAQH/
BAUwAwEB/zAdBgNVHQ4EFgQU7UQZwNPwBovupHu+QucmVMiONnYwDQYJKoZI\nnhvcNAQELBQAD
ggIBAA2ukDL2pkt8RHYZYR4nKM1eVO8lvOMIkPkp165oCOGUAFjvLi5+U1KMtlwH\n6oi6mYtQ
lNeCgN9hCQCTrQ0U5s7B8jeUeLBfnLOic7iPBZM4zy0+sLj7wM+x8uwtLRvM7Kqas6pg\nnghst
O8OEPVeKlh6cdbhTMM1gCIOQ045U8U1mwF10A0Cj7oV+wh93nAbowacyXVKV7cndJZ5t+qnt\n
ozo00F172u1Q8zW/7esUTTHHYPTa8Yec4kjixsU3+wYQ+nVZZjFHKdp2mhZpgq7vmrlR94gjmm
mV\nYjz1VYA211QC//G5Xc7UI2/YRYRKW2XviQzdFKcgyxilJbQN+QHwotL0AMh0jqEqSI512x
PE4iUX\nnfeu+hlsXIFRRk0pTAwvsXcoz7WL9RccvW9xYoIA55vrX/hMUpu09lEpCdNTDd1lzzY
9GvlU47/ro\nnkTLq1lgEIt44w8y8bckzOmoKaT+gyOpyj4xjhi09bTyWnpXgSUyqorkqG5w2gX
jtw+hG4iZZRHUE\n2XWJUc0QhJ1hYMTd+ZciTY6Y5uN/9lu7rs3KSoFrXgvzUeF0K+l+J6fZmU
lO+KWA2yUPHGNIiskz\nZ2s8EIPGrd6ozRaOjfAHN3Gf8qv8QfXBi+wAN10J5U6A7/qxXDgGpR
tK4dw4LTzcqx+QGtVKno7R\ncGzM7vRX+Bi6hG6H\n-----END CERTIFICATE-----\n\"
}"
```

Deleting a certificate using its UUID

```
# The API:
/api/security/certificates/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/certificates/dad2363b-8ac0-11e8-9058-005056b482fc?fields=*" -H "accept: application/hal+json"
```

Signing a new certificate signing request using an existing CA certificate UUID

Once you have created a certificate of type "root_ca", you can use that certificate to act as a local Certificate Authority to sign new certificate signing requests. The following example signs a new certificate signing request using an existing CA certificate UUID. If successful, the API returns a signed certificate.

```
# The API:
/api/security/certificates/{ca.uuid}/sign

# The call:
curl -X POST "https://<mgmt-ip>/api/security/certificates/253add53-8ac9-11e8-9058-005056b482fc/sign" -H "accept: application/json" -H "Content-Type: application/json" -d "{ \"signing_request\": \"-----BEGIN CERTIFICATE REQUEST-----
\nMIICYTCCAUAkCAQAwHDENMAsGA1UEAxMEVEVTVDELMAkGA1UEBhMCVVMwggEiMA0G\nnCSqGSI
b3DQEBAQUAA4IBDwAwggEKAoIBAQCIBCuVfbYHNdOO7vjRQja4JqL2cHqK\ndr1Tj5hz9RVqFK
Z7VPPh8DSP9LoTbYWsvrTkbuD0Wi715MVQCsbkq/mHos+Y5lfqs\nNP5K92fc6EhBzBDYFgZGFn
tZYJjEG5MPerIUE7CfVy7o6sjW0lxeY33pjefObyvP\nBcJkBHg6SFJK/TDLvIYJkonLkJEOJo
TI6++a3I/1bCMfUeuRtLU9ThWlna1kMMYK\n4T16/Bxgm4bha2U2jtosc0Wltnld/capc+eqRV
07WVbMmEOTtop3cv0h3N0S6lbn\nFkd96DXzeGWbSHFHckeCZ9bOHhnVbfEa/efkPLx7ziMC8G
tRHHlwbNk7AgMBAAAG\nADANBgkqhkiG9w0BAQsFAAOCAQEaf+rs1i5PHaOSI2HtTM+Hcv/p71
yzgoLL+aeU\ntB0V4iuoXdqY8oQeWoPI92ci0K08JuSpu6D0DwCK1stfwuGkAA2b0Wr7ZDRonT
Uq\nmJ4j3O47MLysW4Db2LbGws/AuDScIRBJDWHMPHaqsvRbpMx2xQ/V5oagUw5eGGpN\ne4fg
/E2k9mGkpxwkUzT7w1RZirpND4xL+XTzpzzeZqgalpXug4yjiXlI5hpRESZ9/\nAkGJSCWxi15I
ZdxxFVX1Bcm6WpJnnboqkcKeXz95GM6Re+oBy9tlgvwv1Vd5s8uHX+bycFiZp09Wsm8Ev727M
ziZ+0II9nxwkDKsdPvam+KLI9hLQ==\n-----END CERTIFICATE REQUEST-----\n\",
\"hash_function\": \"sha256\"}"

# The response:
{
  "public_certificate": "-----BEGIN CERTIFICATE-----
\nMIIDBzCCAe+gAwIBAgIIFUKQpcqeaUAWDQYJKoZIhvcNAQELBQAWhDENMAsGA1UE\nnAxMEUk
FDWDELMAkGA1UEBhMCVVMwHhcNMTgwNzE4MjAzMTA1WhcNMTkwNzE4MjAz\nnMTA1WjAcMQ0wCw
YDVQQDEWRURVNUMQswCQYDVQQGEWJVUzCCASIwDQYJKoZIhvcN\nnAQEBBQADggEPADCCAQoCgg
EBAKIEK5V9tgc1047u+NFCNrgmovZweop2uVOPmHP1\nnFWoUpntU+HwNI/0uhNthay+tORu4PR
aLvXkxVAKxuSr+Yeiz5jmV+qw0/kr3Z9zo\nnSEHMENgWBkYWellgmMQbkW96shQTSJ9XLujqyN
Y6XF5j femN585vK88FwmQEeDpI\nnUkr9MMu8hgmSicuQkQ4mhMjr75rcj/VsIx9R65G0tT1OFa
WdrWQwxgrhPXR8HGCb\nnHuFrZTa02ixzRaW2eV39xqlz56pFXTtZVsyYQ502indy/SHc3RLqVu
cWR33oNfN4\nnZZtIcUdyR4Jn1s4eGdVt8Rr95+Q8vHvOIwLwa1EceXBucrsCAwEAAnnMEswCQ
YD\nnVR0TBAlwADAdBgNVHQ4EFgQUJMPxjeW1G76TbbD2tXB8dwSpI3MwHwYDVR0jBBgw\nnFoAU
u5aH0mWR4cFoN9i7k96d2op3sPwwDQYJKoZIhvcNAQELBQADggEBAI5ai+Zi\nnFQZUXRTqJCgH
sgBThARneVWQYkYpyAXmTR7QeLfld4ZHL33i4xWCqX3uvW7SFJLe\nnZajT2AVmgIDbaWIHtDtv
qz1BY78PSgUwPH/IyARTEOBeikp6KdwMPraehDIBMAcc\nnANY58wXiTBbsl8UMD6tGecgnzw6s
xlMmadGvrfJeJmgY4zert6NNvgtTPhcZQdLS\nnE0fGzHS6+3ajCCfEEhPNPer9D0e5Me8li9Es
QGENrnJzTci8rzXPuF4bC3gghrK1\nnI1+kmJQ1kLYVUcsntcrIiHmNvtPFJY6stjDgQKS9aDd/
THhPpokPtZoCmE6PDxh6\nnR+dO6C0hcDKHFzA=\n-----END CERTIFICATE-----\n"
}
```


Retrieve security certificates

GET /security/certificates

Introduced In: 9.6

Retrieves security certificates.

Related ONTAP commands

- `security certificate show`

Parameters

Name	Type	In	Required	Description
public_certificate	string	query	False	Filter by public_certificate <ul style="list-style-type: none">• Introduced in: 9.8
hash_function	string	query	False	Filter by hash_function
common_name	string	query	False	Filter by common_name
authority_key_identifier	string	query	False	Filter by authority_key_identifier <ul style="list-style-type: none">• Introduced in: 9.8
private_key	string	query	False	Filter by private_key <ul style="list-style-type: none">• Introduced in: 9.8
type	string	query	False	Filter by type
uuid	string	query	False	Filter by uuid <ul style="list-style-type: none">• Introduced in: 9.8
name	string	query	False	Filter by name <ul style="list-style-type: none">• Introduced in: 9.8

Name	Type	In	Required	Description
subject_key_identifier	string	query	False	Filter by subject_key_identifier • Introduced in: 9.8
serial_number	string	query	False	Filter by serial_number
key_size	integer	query	False	Filter by key_size
ca	string	query	False	Filter by ca
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
expiry_time	string	query	False	Filter by expiry_time
intermediate_certificates	string	query	False	Filter by intermediate_certificates • Introduced in: 9.8
scope	string	query	False	Filter by scope
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[security_certificate]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authority_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D7",
      "ca": "string",
      "common_name": "test.domain.com",
      "expiry_time": "string",
      "hash_function": "string",
      "intermediate_certificates": [
        "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAwHDEnMAkGA1UE
AxMEVEVTVDELMAkGA1UEBhMCVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBBYEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKA FMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA
vDovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklkBtVBDTmLnrc -----END CERTIFICATE-----"
      ],
      "name": "cert1",
      "private_key": "-----BEGIN PRIVATE KEY-----
MIIBVAIBADANBgkqhkiG9w0BAQEFAASCAT4wggE6AgEAAKEAu1/a8f3G47cZ6pel
Hd3aONMNkGJ8vSCH5QjicuDm92VtVwkAACEjIoZSLYLJvPD+odL+lFzVQSmkneW7
VCGqYQIDAQABAKAcfNpg6GCQxoneLOghvlUrRotNZGvqpUOEAvHK3X7AJhZ5SU4V
an36qvsAt5ghFMVM2iGvGaXbj0dAd+Jg64pxAiEA32Eh9mPtFSmZhTIUMeGcPmPk
qIYCEuP8a/ZLmI9s4TsCIQDWvLQuvjSVfwPhi0TFab5wqAET8X5LBFqtGX5QlUep
EwIgFnqM02Gc4wtLoqa2d4qPkYu13+uUW9hLd4XSd6i/OS8CIQDT3elU+Rt+qIwW
u0cFrVvNYSV3HNzDfS9N/IoxTagfewIgpVxAdE5c2EWbhCUkhN+ZCf38AKewK9TW
```

```

lQcDy4L+f14= -----END PRIVATE KEY-----",
    "public_certificate": "-----BEGIN CERTIFICATE-----
MIIBuzCCAwwgAwIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAwhDENMAsgA1UE
AxMEVEVTVDELMAkGA1UEBhMCVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBBYEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKAQFMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA
vDovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklnkBTvBDTmLnrc -----END CERTIFICATE-----",
    "scope": "string",
    "serial_number": "string",
    "subject_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D8",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svml",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "type": "string",
    "uuid": "string"
  }
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

security_certificate

Name	Type	Description
_links	_links	
authority_key_identifier	string	Provides the key identifier of the issuing CA certificate that signed the SSL certificate.
ca	string	Certificate authority
common_name	string	FQDN or custom common name. Provide on POST when creating a self-signed certificate.

Name	Type	Description
expiry_time	string	Certificate expiration time. Can be provided on POST if creating self-signed certificate. The expiration time range is between 1 day to 10 years.
hash_function	string	Hashing function. Can be provided on POST when creating a self-signed certificate. Hash functions md5 and sha1 are not allowed on POST.
intermediate_certificates	array[string]	Chain of intermediate Certificates in PEM format. Only valid in POST when installing a certificate.
key_size	integer	Key size of requested Certificate in bits. One of 512, 1024, 1536, 2048, 3072. Can be provided on POST if creating self-signed certificate. Key size of 512 is not allowed on POST.
name	string	Certificate name. If not provided in POST, a unique name specific to the SVM is automatically generated.
private_key	string	Private key Certificate in PEM format. Only valid for create when installing a CA-signed certificate. This is not audited.
public_certificate	string	Public key Certificate in PEM format. If this is not provided in POST, a self-signed certificate is created.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
serial_number	string	Serial number of certificate.
subject_key_identifier	string	Provides the key identifier used to identify the public key in the SSL certificate.

Name	Type	Description
svm	svm	
type	string	<p>Type of Certificate. The following types are supported:</p> <ul style="list-style-type: none"> • client - a certificate and its private key used by an SSL client in ONTAP. • server - a certificate and its private key used by an SSL server in ONTAP. • client_ca - a Certificate Authority certificate used by an SSL server in ONTAP to verify an SSL client certificate. • server_ca - a Certificate Authority certificate used by an SSL client in ONTAP to verify an SSL server certificate. • root_ca - a self-signed certificate used by ONTAP to sign other certificates by acting as a Certificate Authority. • enum: ["client", "server", "client_ca", "server_ca", "root_ca"] • Introduced in: 9.6
uuid	string	Unique ID that identifies a certificate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create or install security certificates

POST /security/certificates

Introduced In: 9.6

Creates or installs a certificate.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create or install the certificate.
- `common_name` - Common name of the certificate. Required when creating a certificate.
- `type` - Type of certificate.
- `public_certificate` - Public key certificate in PEM format. Required when installing a certificate.
- `private_key` - Private key certificate in PEM format. Required when installing a CA-signed certificate.

Recommended optional properties

- `expiry_time` - Certificate expiration time. Specifying an expiration time is recommended when creating a certificate.
- `key_size` - Key size of the certificate in bits. Specifying a strong key size is recommended when creating a certificate.
- `name` - Unique certificate name per SVM. If one is not provided, it is automatically generated.

Default property values

If not specified in POST, the following default property values are assigned:

- `key_size` - *2048*
- `expiry_time` - *P365DT*
- `hash_function` - *sha256*

Related ONTAP commands

- `security certificate create`
- `security certificate install`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
authority_key_identifier	string	Provides the key identifier of the issuing CA certificate that signed the SSL certificate.
ca	string	Certificate authority
common_name	string	FQDN or custom common name. Provide on POST when creating a self-signed certificate.
expiry_time	string	Certificate expiration time. Can be provided on POST if creating self-signed certificate. The expiration time range is between 1 day to 10 years.
hash_function	string	Hashing function. Can be provided on POST when creating a self-signed certificate. Hash functions md5 and sha1 are not allowed on POST.
intermediate_certificates	array[string]	Chain of intermediate Certificates in PEM format. Only valid in POST when installing a certificate.
key_size	integer	Key size of requested Certificate in bits. One of 512, 1024, 1536, 2048, 3072. Can be provided on POST if creating self-signed certificate. Key size of 512 is not allowed on POST.

Name	Type	Description
name	string	Certificate name. If not provided in POST, a unique name specific to the SVM is automatically generated.
private_key	string	Private key Certificate in PEM format. Only valid for create when installing a CA-signed certificate. This is not audited.
public_certificate	string	Public key Certificate in PEM format. If this is not provided in POST, a self-signed certificate is created.
serial_number	string	Serial number of certificate.
subject_key_identifier	string	Provides the key identifier used to identify the public key in the SSL certificate.
svm	svm	

Name	Type	Description
type	string	<p>Type of Certificate. The following types are supported:</p> <ul style="list-style-type: none"> • client - a certificate and its private key used by an SSL client in ONTAP. • server - a certificate and its private key used by an SSL server in ONTAP. • client_ca - a Certificate Authority certificate used by an SSL server in ONTAP to verify an SSL client certificate. • server_ca - a Certificate Authority certificate used by an SSL client in ONTAP to verify an SSL server certificate. • root_ca - a self-signed certificate used by ONTAP to sign other certificates by acting as a Certificate Authority. • enum: ["client", "server", "client_ca", "server_ca", "root_ca"] • Introduced in: 9.6
uuid	string	Unique ID that identifies a certificate.

Example request

```
{
  "authority_key_identifier":
    "26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D7",
  "ca": "string",
  "common_name": "test.domain.com",
  "expiry_time": "string",
  "hash_function": "string",
  "intermediate_certificates": [
    "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAWhDENMAsGA1UE
AxMEVEVTVDELMAkGA1UEBhMCVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBByEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKA FMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA
vDovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklknbTbVBDTmLnrc -----END CERTIFICATE-----"
  ],
  "name": "cert1",
  "private_key": "-----BEGIN PRIVATE KEY-----
MIIBVAIBADANBgkqhkiG9w0BAQEFAASCAT4wggE6AgEAAkEAu1/a8f3G47cZ6pel
Hd3aONMNkGJ8vSCH5QjicuDm92VtVwkAACEjIoZSLYlJvPD+odL+lFzVQSmkneW7
VCGqYQIDAQABAkAcfNpg6GCQxoneLOghvlUrRotNZGvqpUOEAvHK3X7AJhz5SU4V
an36qvsAt5ghFMVM2iGvGaXbj0dAd+Jg64pxAiEA32Eh9mPtFSmZhTIUMeGcPmPk
qIYCEuP8a/ZLmI9s4TsCIQDWvLQuvjSVfwPhi0TFab5wqAET8X5LBFqtGX5QlUep
EwIgFnqM02Gc4wtLoqa2d4qPkYu13+uUW9hLd4XSd6i/OS8CIQDT3elU+Rt+qIwW
u0cFrVvNYSV3HNzDfS9N/IoxTagfewIgpVxAdE5c2EWbhCUkhN+ZCf38AKewK9TW
lQcDy4L+f14= -----END PRIVATE KEY-----",
  "public_certificate": "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAWhDENMAsGA1UE
AxMEVEVTVDELMAkGA1UEBhMCVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBByEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKA FMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA
vDovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklknbTbVBDTmLnrc -----END CERTIFICATE-----",
  "serial_number": "string",
  "subject_key_identifier":
    "26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D8",
}
```

```
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"type": "string",
"uuid": "string"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[security_certificate]	

Example response

```
{
  "records": [
    {
      "authority_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D7",
      "ca": "string",
      "common_name": "test.domain.com",
      "expiry_time": "string",
      "hash_function": "string",
      "intermediate_certificates": [
        "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAwhDENMAsGA1UE
AxMEVEVTVDELMAkGA1UEBhMCVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBByEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKA FMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA
vDovYeyGNnknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklnkBTvBDTmLnrc -----END CERTIFICATE-----"
      ],
      "name": "cert1",
      "private_key": "-----BEGIN PRIVATE KEY-----
MIIBVAIBADANBgkqhkiG9w0BAQEFAASCAT4wggE6AgEAAkEAu1/a8f3G47cZ6pel
Hd3aONMNkGJ8vSCH5QjicuDm92VtVwKAACEjIoZSLYlJvPD+odL+lFzVQSmkneW7
VCGqYQIDAQABAkAcfNpg6GCQxoneLOghv1UrRotNZGvqpUOEAvHK3X7AJhz5SU4V
an36qvsAt5ghFMVM2iGvGaXbj0dAd+Jg64pxAiEA32Eh9mPtFSmZhTIUMeGcPmPk
qIYCEuP8a/ZLmI9s4TsCIQDWvLQuvJsvfwPhi0TFab5wqAET8X5LBFqtGX5QlUep
EwIgFnqM02Gc4wtLoqa2d4qPkYu13+uUW9hLd4XSd6i/OS8CIQDT3elU+Rt+qIwW
u0cFrVvNYSV3HNzDfS9N/IoxTagfewIgpVxAdE5c2EWbhCUkhN+ZCf38AKewK9TW
lQcDy4L+f14= -----END PRIVATE KEY-----",
      "public_certificate": "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAwhDENMAsGA1UE
AxMEVEVTVDELMAkGA1UEBhMCVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVR0TBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBByEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKA FMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA
vDovYeyGNnknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklnkBTvBDTmLnrc -----END CERTIFICATE-----",
      "serial_number": "string",
    }
  ]
}
```



```

    "subject_key_identifier":
      "26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D8",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "uuid": "string"
    }
  ]
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
3735645	Cannot specify a value for serial. It is generated automatically.
3735622	The certificate type is not supported.
3735664	The specified key size is not supported in FIPS mode.
3735665	The specified hash function is not supported in FIPS mode.
3735553	Failed to create self-signed Certificate.
3735646	Failed to store the certificates.
3735693	The certificate installation failed as private key was empty.
3735618	Cannot accept private key for server_ca or client_ca.
52363365	Failed to allocate memory.
52559975	Failed to read the certificate due to incorrect formatting.
52363366	Unsupported key type.
52560123	Failed to read the key due to incorrect formatting.
52559972	The certificates start date is later than the current date.
52559976	The certificate and private key do not match.
52559973	The certificate has expired.

Error Code	Description
52363366	Logic error: use of a dead object.
3735696	Intermediate certificates are not supported with client_ca and server_ca type certificates.
52559974	The certificate is not supported in FIPS mode.
3735676	Cannot continue the installation without a value for the common name. Since the subject field in the certificate is empty, the field "common_name" must have a value to continue with the installation.
3735558	Failed to extract information about Common Name from the certificate.
3735588	The common name (CN) extracted from the certificate is not valid.
3735632	Failed to extract Certificate Authority Information from the certificate.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

security_certificate

Name	Type	Description
authority_key_identifier	string	Provides the key identifier of the issuing CA certificate that signed the SSL certificate.
ca	string	Certificate authority
common_name	string	FQDN or custom common name. Provide on POST when creating a self-signed certificate.
expiry_time	string	Certificate expiration time. Can be provided on POST if creating self-signed certificate. The expiration time range is between 1 day to 10 years.
hash_function	string	Hashing function. Can be provided on POST when creating a self-signed certificate. Hash functions md5 and sha1 are not allowed on POST.
intermediate_certificates	array[string]	Chain of intermediate Certificates in PEM format. Only valid in POST when installing a certificate.

Name	Type	Description
key_size	integer	Key size of requested Certificate in bits. One of 512, 1024, 1536, 2048, 3072. Can be provided on POST if creating self-signed certificate. Key size of 512 is not allowed on POST.
name	string	Certificate name. If not provided in POST, a unique name specific to the SVM is automatically generated.
private_key	string	Private key Certificate in PEM format. Only valid for create when installing a CA-signed certificate. This is not audited.
public_certificate	string	Public key Certificate in PEM format. If this is not provided in POST, a self-signed certificate is created.
serial_number	string	Serial number of certificate.
subject_key_identifier	string	Provides the key identifier used to identify the public key in the SSL certificate.
svm	svm	

Name	Type	Description
type	string	<p>Type of Certificate. The following types are supported:</p> <ul style="list-style-type: none"> • client - a certificate and its private key used by an SSL client in ONTAP. • server - a certificate and its private key used by an SSL server in ONTAP. • client_ca - a Certificate Authority certificate used by an SSL server in ONTAP to verify an SSL client certificate. • server_ca - a Certificate Authority certificate used by an SSL client in ONTAP to verify an SSL server certificate. • root_ca - a self-signed certificate used by ONTAP to sign other certificates by acting as a Certificate Authority. • enum: ["client", "server", "client_ca", "server_ca", "root_ca"] • Introduced in: 9.6
uuid	string	Unique ID that identifies a certificate.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Sign security certificates

POST /security/certificates/{ca.uuid}/sign

Introduced In: 9.6

Signs a certificate.

Required properties

- `signing_request` - Certificate signing request to be signed by the given certificate authority.

Recommended optional properties

- `expiry_time` - Certificate expiration time. Specifying an expiration time for a signed certificate is recommended.
- `hash_function` - Hashing function. Specifying a strong hashing function is recommended when signing a certificate.

Default property values

If not specified in POST, the following default property values are assigned:

- `expiry_time` - *P365DT*
- `hash_function` - *sha256*

Related ONTAP commands

- `security certificate sign`
This API is used to sign a certificate request using a pre-existing self-signed root certificate. The self-signed root certificate acts as a certificate authority within its scope and maintains the records of its signed certificates.

The root certificate can be created for a given SVM or for the cluster using [POST security/certificates].

Parameters

Name	Type	In	Required	Description
ca.uuid	string	path	True	UUID of the existing certificate authority certificate
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
expiry_time	string	Certificate expiration time. The allowed expiration time range is between 1 day to 10 years.
hash_function	string	Hashing function
signing_request	string	Certificate signing request to be signed by the given certificate authority. Request should be in X509 PEM format.

Example request

```
{
  "expiry_time": "string",
  "hash_function": "string",
  "signing_request": "'-----BEGIN CERTIFICATE REQUEST-----
MIICYDCCAUGCAQAwGzEMMAoGA1UEAxMDQUJDMQswCQYDVQQGEwJVUzCCASIwDQYJ
KoZIhvcNAQEBBQADggEPADCCAQoCggEBAPF+82SlqT3Vyu3Jx4IAwHcO5EGwLOxy
zQ6KNjz71Fca0n1/A1CbCPyOsSupGVObvdWxX7xLVMJ2SXB7h43GCqYyX6FXJO4F
HOpmLvB+jxdeiW7SDbiZyLUlsvA+oRO/uNlcug773QZdKLjJD64erZZMRUNbUJB8
bARxAUi0FPvgTraSQ0UW5sRLiGKeAyKA4wekYe1VgjHRTBizFbD4dI3njfva/2B1
jf+kkulgcLJTUJNtkgeimqMKyraYuleYcYk2K+C//0NuNOuPbDfTXCM7O61vik09
Szi8nLN7OXE9KoAA93U/BCpSfpl8XIb4cGnEr8hgVHOotZSo+KZBFxMCAwEAAaAA
MA0GCSqGSIB3DQEBChUAA4IBAQc2vFYpvgsFrm5GnPx8tOBD1xsTyYjbWJMD8hAF
lFrvF9Sw9QGctDyacxkwgJhQx8l8JiIS5GOY6WWLB19FMkLQNAhDL9xF3WF7vfYq
RKgrz3bd/Vg96fsRZNYIPLGmoEaqLOh3FOCGc2VbdsR9PwOn3fwthxkIRd6ds6/q
jc5cpSmVsCOgu+OKcpRXikYDbkWXfTZ1AhSfn6njBYFdZ9+PNAu/0JRQh5bX60nO
5heniTcaJLwUZP/CQ8nxHY0Wqy+lrAtM33d5cVmhUlBXQSIru/0ZkA/b9fK5Zv8E
ZMADYUoEvIG59VxhyCi8lzYf+Mxl8qBSF+ZdC4yWhzDqZtM9 -----END CERTIFICATE
REQUEST-----'"
}
```

Response

Status: 200, Ok

Name	Type	Description
public_certificate	string	CA signed public key Certificate

Example response

```
{
  "public_certificate": "string"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
3735628	Failed to use CA certificate for signing.
3735665	The specified hash function is not supported in FIPS mode.
52559974	The certificate is not supported in FIPS mode.
3735626	Failed to generate signed Certificate.
3735558	Failed to extract information about Common Name from the certificate.
3735588	The common name (CN) extracted from the certificate is not valid.
3735632	Failed to extract Certificate Authority Information from the certificate.
3735629	Failed to sign the certificate because Common Name of signing certificate and Common Name of CA certificate are same.
3735630	Failed to sign the certificate because expiry date of signing certificate exceeds the expiry date of CA certificate.

Definitions

See Definitions

security_certificate_sign

Name	Type	Description
expiry_time	string	Certificate expiration time. The allowed expiration time range is between 1 day to 10 years.
hash_function	string	Hashing function
signing_request	string	Certificate signing request to be signed by the given certificate authority. Request should be in X509 PEM format.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete security certificates

DELETE /security/certificates/{uuid}

Introduced In: 9.6

Deletes a security certificate.

Related ONTAP commands

- `security certificate delete`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Certificate UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
3735644	Cannot delete server-chain certificate. Reason: There is a corresponding server certificate for it.
3735679	Cannot delete pre-installed server_ca certificates through REST. Use CLI or ZAPI.
3735650	Deleting this client_ca certificate directly is not supported. Delete the corresponding root-ca certificate using type <code>root_ca</code> to delete the root, client, and server certificates.
3735627	Deleting this server_ca certificate directly is not supported. Delete the corresponding root-ca certificate using type <code>root_ca</code> to delete the root, client, and server certificates.
3735589	Cannot delete certificate.
3735590	Cannot delete certificate. Failed to remove SSL configuration for the certificate.
3735683	Cannot remove this certificate while external key manager is configured.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve security certificates

GET /security/certificates/{uuid}

Introduced In: 9.6

Retrieves security certificates.

Related ONTAP commands

- `security certificate show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Certificate UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
authority_key_identifier	string	Provides the key identifier of the issuing CA certificate that signed the SSL certificate.
ca	string	Certificate authority
common_name	string	FQDN or custom common name. Provide on POST when creating a self-signed certificate.
expiry_time	string	Certificate expiration time. Can be provided on POST if creating self-signed certificate. The expiration time range is between 1 day to 10 years.
hash_function	string	Hashing function. Can be provided on POST when creating a self-signed certificate. Hash functions md5 and sha1 are not allowed on POST.

Name	Type	Description
intermediate_certificates	array[string]	Chain of intermediate Certificates in PEM format. Only valid in POST when installing a certificate.
key_size	integer	Key size of requested Certificate in bits. One of 512, 1024, 1536, 2048, 3072. Can be provided on POST if creating self-signed certificate. Key size of 512 is not allowed on POST.
name	string	Certificate name. If not provided in POST, a unique name specific to the SVM is automatically generated.
private_key	string	Private key Certificate in PEM format. Only valid for create when installing a CA-signed certificate. This is not audited.
public_certificate	string	Public key Certificate in PEM format. If this is not provided in POST, a self-signed certificate is created.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
serial_number	string	Serial number of certificate.
subject_key_identifier	string	Provides the key identifier used to identify the public key in the SSL certificate.
svm	svm	

Name	Type	Description
type	string	<p>Type of Certificate. The following types are supported:</p> <ul style="list-style-type: none"> • client - a certificate and its private key used by an SSL client in ONTAP. • server - a certificate and its private key used by an SSL server in ONTAP. • client_ca - a Certificate Authority certificate used by an SSL server in ONTAP to verify an SSL client certificate. • server_ca - a Certificate Authority certificate used by an SSL client in ONTAP to verify an SSL server certificate. • root_ca - a self-signed certificate used by ONTAP to sign other certificates by acting as a Certificate Authority. • enum: ["client", "server", "client_ca", "server_ca", "root_ca"] • Introduced in: 9.6
uuid	string	Unique ID that identifies a certificate.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authority_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D7",
  "ca": "string",
  "common_name": "test.domain.com",
  "expiry_time": "string",
  "hash_function": "string",
  "intermediate_certificates": [
    "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAWhDENMAsGA1UE
AxMEVEVTVDELMAkGA1UEBhMCVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVROTBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBByEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKA FMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA v
DovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklkBtVBDTmLnrC -----END CERTIFICATE-----"
  ],
  "name": "cert1",
  "private_key": "-----BEGIN PRIVATE KEY-----
MIIBVAIBADANBgkqhkiG9w0BAQEFAASCAT4wggE6AgEAAkEAu1/a8f3G47cZ6pel
Hd3aONMNkGJ8vSCH5QjicuDm92VtVwkaACEjIoZSLYlJvPD+odL+lfzVQSmkneW7
VCGqYQIDAQABAkAcfNpg6GCQxoneLOghvlUrRotNZGvqpUOEAvHK3X7AJhZ5SU4V
an36qvsAt5ghFMVM2iGvGaXbj0dAd+Jg64pxAiEA32Eh9mPtFsmZhTIUMeGcPmPk
qIYCEuP8a/ZLmI9s4TsCIQDWvLQuvjSVfwPhi0TFab5wqAET8X5LBFqtGX5QlUep
EwIgfFnqM02Gc4wtLoqa2d4qPkYu13+uUW9hLd4XSd6i/OS8CIQDT3elU+Rt+qIwW
u0cFrVvNYSV3HNzDfS9N/IoxTagfewIgpVXADe5c2EWbhCUkhN+ZCf38AKewK9TW
lQcDy4L+f14= -----END PRIVATE KEY-----",
  "public_certificate": "-----BEGIN CERTIFICATE-----
MIIBuzCCAWWgAwIBAgIIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQAWhDENMAsGA1UE
AxMEVEVTVDELMAkGA1UEBhMCVVMwHhcNMTgwNjA4MTgwOTAxWhcNMTkwNjA4MTgw
OTAxWjAcMQ0wCwYDVQQDEwRURVNUMQswCQYDVQQGEwJVUzBcMA0GCSqGSIb3DQEB
AQUAA0sAMEgCQQDaPvbqUJJFJ6NNTyK3Yb+ytSjJ9aa3yUmYTD9uMiP+6ycjxHWB
e8u9z6yCHsW03ync+dnhE5c5z8wuDAY0fv15AgMBAAGjgYowgYcwDAYDVROTBAUw
AwEB/zALBgNVHQ8EBAMCAQYwHQYDVR0OBByEFMJ7Ev/o/3+YNzYh5XNlqqjnw4zm
MEsGA1UdIwREMEKA FMJ7Ev/o/3+YNzYh5XNlqqjnw4zmoSCkHjAcMQ0wCwYDVQQD
EwRURVNUMQswCQYDVQQGEwJVU4IIFTZBrqZwUUMwDQYJKoZIhvcNAQELBQADQQA v
```



```
DovYeyGNknjGI+TVNX6nDbyzf7zUPqnri0KuvObEeybrbPW45sgsnT5dyeE/32U
9Yr6lklnkBtVBDTmLnrc -----END CERTIFICATE-----",
  "scope": "string",
  "serial_number": "string",
  "subject_key_identifier":
"26:1F:C5:53:5B:D7:9E:E2:37:74:F4:F4:06:09:03:3D:EB:41:75:D8",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View and update IPsec configuration

Security IPsec endpoint overview

Overview

The following operations are supported:

- GET to retrieve the IPsec status: GET security/ipsec
- Patch to update IPsec status: PATCH security/ipsec

Retrieve an IPsec configuration

GET /security/ipsec

Introduced In: 9.8

Retrieves IPsec configuration via REST APIs.

Related ONTAP commands

- 'security ipsec config show'

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	self_link	
enabled	boolean	Indicates whether or not IPsec is enabled.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update an IPsec configuration

PATCH /security/ipsec

Introduced In: 9.8

Updates IPsec configuration via REST APIs.

Related ONTAP commands

- 'security ipsec config modify'

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
enabled	boolean	Indicates whether or not IPsec is enabled.

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66256898	Internal error. Failed to enable IPsec.
66256899	Internal error. Failed to disable IPsec.
66257199	IPsec is not supported in the current cluster version.

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

ipsec

Manages IPsec configuration via REST APIs.

Name	Type	Description
enabled	boolean	Indicates whether or not IPsec is enabled.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage IPsec policies

Security IPsec policies endpoint overview

Overview

The following operations are supported:

- Collection Get: GET security/ipsec/policies

- Creation Post: POST security/ipsec/policies
- Instance Get: GET security/ipsec/policies/uuid
- Instance Patch: PATCH security/ipsec/policies/uuid
- Instance Delete: DELETE security/ipsec/policies/uuid

Retrieve IPsec policies

GET /security/ipsec/policies

Introduced In: 9.8

Retrieves the collection of IPsec policies.

Related ONTAP commands

- security ipsec policy show

Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
name	string	query	False	Filter by name
protocol	string	query	False	Filter by protocol
scope	string	query	False	Filter by scope
remote_endpoint.net mask	string	query	False	Filter by remote_endpoint.net mask
remote_endpoint.fa mily	string	query	False	Filter by remote_endpoint.fa mily
remote_endpoint.ad dress	string	query	False	Filter by remote_endpoint.ad dress
enabled	boolean	query	False	Filter by enabled
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name

Name	Type	In	Required	Description
local_endpoint.netmask	string	query	False	Filter by local_endpoint.netmask
local_endpoint.family	string	query	False	Filter by local_endpoint.family
local_endpoint.address	string	query	False	Filter by local_endpoint.address
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
error	error	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "num_records": "1",
  "records": [
    {
      "local_endpoint": {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24"
      },
      "name": "string",
      "protocol": "17",
      "remote_endpoint": {
        "address": "10.10.10.7",
        "family": "string",
        "netmask": "24"
      },
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

```
    },
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

local_endpoint

Local endpoint for the IPsec policy.

Name	Type	Description
address	string	IPv4 or IPv6 address
family	string	IPv4 or IPv6

Name	Type	Description
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

remote_endpoint

Remote endpoint for the IPsec policy.

Name	Type	Description
address	string	IPv4 or IPv6 address
family	string	IPv4 or IPv6
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

_links

Name	Type	Description
self	href	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

records

IPsec policy object.

Name	Type	Description
enabled	boolean	Indicates whether or not the policy is enabled.

Name	Type	Description
local_endpoint	local_endpoint	Local endpoint for the IPsec policy.
name	string	IPsec policy name.
protocol	string	Lower layer protocol to be covered by the IPsec policy.
remote_endpoint	remote_endpoint	Remote endpoint for the IPsec policy.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	Unique identifier of the IPsec policy.

Create an IPsec policy

POST /security/ipsec/policies

Introduced In: 9.8

Creates an IPsec policy.

Related ONTAP commands

- `security ipsec policy create`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
action	string	Action for the IPsec policy.
enabled	boolean	Indicates whether or not the policy is enabled.
local_endpoint	local_endpoint	Local endpoint for the IPsec policy.
name	string	IPsec policy name.
protocol	string	Lower layer protocol to be covered by the IPsec policy.
remote_endpoint	remote_endpoint	Remote endpoint for the IPsec policy.
secret_key	string	Pre-shared key for IKE negotiation.
svm	svm	
uuid	string	Unique identifier of the IPsec policy.

Example request

```
{
  "action": "string",
  "local_endpoint": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "name": "string",
  "protocol": "17",
  "remote_endpoint": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "secret_key": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "num_records": "1",
  "records": [
    {
      "action": "string",
      "local_endpoint": {
        "address": "10.10.10.7",
        "netmask": "24"
      },
      "name": "string",
      "protocol": "17",
      "remote_endpoint": {
        "address": "10.10.10.7",
        "netmask": "24"
      },
      "secret_key": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66257099	Only one protocol can be specified.
66257100	Only one local port can be specified.
66257101	Only one remote port can be specified.
66257103	IPsec policy with same name already exists in this SVM.
66257107	The specified pre-shared key is not a valid hexadecimal string.
66257109	The specified pre-shared key is not a valid Base64 encoded binary string.
66257110	Failed to a create policy sequencing value.
66257111	The IPsec policy with action ESP TRANSPORT provides packet protection and requires a secret key for authentication.
66257112	The IPsec policy with the action specified does not provide packet protection and the authentication method provided for the policy will be ignored.
66257113	Only one local IP subnet can be specified.
66257114	Only one remote IP subnet can be specified.
66257115	Port ranges containing more than one port are not supported.
66257117	IPsec is not supported on the SVM specified in the policy, IPsec is supported on data SVMs only.
66257120	The subnet selector must be a host address (An IPv4 address with a 32-bit netmask or an IPv6 address with a 128-bit netmask).
66257121	The maximum limit of IPsec Policies has reached for the specified SVM.
66257125	The local_endpoint.address must be specified with local_endpoint.netmask.
66257126	The remote_endpoint.address must be specified with remote_endpoint.netmask.

Definitions

See Definitions

local_endpoint

Local endpoint for the IPsec policy.

Name	Type	Description
address	string	IPv4 or IPv6 address
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

remote_endpoint

Remote endpoint for the IPsec policy.

Name	Type	Description
address	string	IPv4 or IPv6 address
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipsec_policy

IPsec policy object.

Name	Type	Description
action	string	Action for the IPsec policy.
enabled	boolean	Indicates whether or not the policy is enabled.
local_endpoint	local_endpoint	Local endpoint for the IPsec policy.
name	string	IPsec policy name.
protocol	string	Lower layer protocol to be covered by the IPsec policy.
remote_endpoint	remote_endpoint	Remote endpoint for the IPsec policy.
secret_key	string	Pre-shared key for IKE negotiation.
svm	svm	
uuid	string	Unique identifier of the IPsec policy.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

records

IPsec policy object.

Name	Type	Description
action	string	Action for the IPsec policy.
enabled	boolean	Indicates whether or not the policy is enabled.
local_endpoint	local_endpoint	Local endpoint for the IPsec policy.
name	string	IPsec policy name.
protocol	string	Lower layer protocol to be covered by the IPsec policy.
remote_endpoint	remote_endpoint	Remote endpoint for the IPsec policy.
secret_key	string	Pre-shared key for IKE negotiation.
svm	svm	SVM, applies only to SVM-scoped objects.
uuid	string	Unique identifier of the IPsec policy.

Delete an IPsec policy

```
DELETE /security/ipsec/policies/{uuid}
```

Introduced In: 9.8

Deletes a specific IPsec policy.

Related ONTAP commands

- `security ipsec policy delete`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	IPsec policy UUID

Response

```
Status: 200, Ok
```

Error

```
Status: Default
```

ONTAP Error Response Codes

Error Code	Description
66257096	Internal error. Failed to purge connections associated with the IPsec policy.
66257116	IPsec policy with the specified UUID was not found.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve an IPsec policy

GET /security/ipsec/policies/{uuid}

Introduced In: 9.8

Retrieves a specific IPsec policy.

Related ONTAP commands

- `security ipsec policy show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	IPsec policy UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
enabled	boolean	Indicates whether or not the policy is enabled.
local_endpoint	local_endpoint	Local endpoint for the IPsec policy.
name	string	IPsec policy name.
protocol	string	Lower layer protocol to be covered by the IPsec policy.
remote_endpoint	remote_endpoint	Remote endpoint for the IPsec policy.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
svm	svm	
uuid	string	Unique identifier of the IPsec policy.

Example response

```
{
  "local_endpoint": {
    "address": "10.10.10.7",
    "family": "string",
    "netmask": "24"
  },
  "name": "string",
  "protocol": "17",
  "remote_endpoint": {
    "address": "10.10.10.7",
    "family": "string",
    "netmask": "24"
  },
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description	
66257116	IPsec policy with the specified UUID was not found.	

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

local_endpoint

Local endpoint for the IPsec policy.

Name	Type	Description
address	string	IPv4 or IPv6 address
family	string	IPv4 or IPv6
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

remote_endpoint

Remote endpoint for the IPsec policy.

Name	Type	Description
address	string	IPv4 or IPv6 address
family	string	IPv4 or IPv6
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update an IPsec policy

PATCH /security/ipsec/policies/{uuid}

Introduced In: 9.8

Updates a specific IPsec policy.

Related ONTAP commands

- `security ipsec policy modify`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	IPsec policy UUID

Request Body

Name	Type	Description
enabled	boolean	Indicates whether or not the policy is enabled.
local_endpoint	local_endpoint	Local endpoint for the IPsec policy.
name	string	IPsec policy name.
protocol	string	Lower layer protocol to be covered by the IPsec policy.
remote_endpoint	remote_endpoint	Remote endpoint for the IPsec policy.
svm	svm	
uuid	string	Unique identifier of the IPsec policy.

Example request

```
{
  "local_endpoint": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "name": "string",
  "protocol": "17",
  "remote_endpoint": {
    "address": "10.10.10.7",
    "netmask": "24"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66257097	Internal error. Failed to update the IPsec policy.
66257099	Only one protocol can be specified.
66257100	Only one local port can be specified.
66257101	Only one remote port can be specified.
66257110	Failed to create a policy sequencing value.
66257113	Only one local IP subnet can be specified.
66257114	Only one remote IP subnet can be specified.
66257115	Port ranges containing more than one port are not supported.
66257116	IPsec policy with the specified UUID was not found.
66257120	The subnet selector must be a host address (An IPv4 address with a 32-bit netmask or an IPv6 address with a 128-bit netmask).

Definitions

See Definitions

local_endpoint

Local endpoint for the IPsec policy.

Name	Type	Description
address	string	IPv4 or IPv6 address
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

remote_endpoint

Remote endpoint for the IPsec policy.

Name	Type	Description
address	string	IPv4 or IPv6 address
netmask	string	Input as netmask length (16) or IPv4 mask (255.255.0.0). For IPv6, you must set the netmask length. The default value is 64. Output is always netmask length.

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

ipsec_policy

IPsec policy object.

Name	Type	Description
enabled	boolean	Indicates whether or not the policy is enabled.
local_endpoint	local_endpoint	Local endpoint for the IPsec policy.
name	string	IPsec policy name.
protocol	string	Lower layer protocol to be covered by the IPsec policy.
remote_endpoint	remote_endpoint	Remote endpoint for the IPsec policy.
svm	svm	
uuid	string	Unique identifier of the IPsec policy.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View IPsec security associations

Security IPsec security-associations endpoint overview

Overview

- Collection Get: GET security/ipsec/security-associations
- Instance Get: GET security/ipsec/security-associations/uuid

Retrieve IPsec and IKE security associations

GET /security/ipsec/security-associations

Introduced In: 9.8

Retrieves the IPsec and IKE (Internet Key Exchange) security associations.

Related ONTAP commands

- `security ipsec show-ipseca`
- `security ipsec show-ikesa`

Parameters

Name	Type	In	Required	Description
local_address	string	query	False	Filter by local_address
lifetime	integer	query	False	Filter by lifetime
remote_address	string	query	False	Filter by remote_address
cipher_suite	string	query	False	Filter by cipher_suite
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
type	string	query	False	Filter by type
ike.version	string	query	False	Filter by ike.version
ike.authentication	string	query	False	Filter by ike.authentication

Name	Type	In	Required	Description
ike.initiator_security_parameter_index	integer	query	False	Filter by ike.initiator_security_parameter_index
ike.state	string	query	False	Filter by ike.state
ike.responder_security_parameter_index	integer	query	False	Filter by ike.responder_security_parameter_index
uuid	string	query	False	Filter by uuid
ipsec.state	string	query	False	Filter by ipsec.state
ipsec.outbound.bytes	integer	query	False	Filter by ipsec.outbound.bytes
ipsec.outbound.packets	integer	query	False	Filter by ipsec.outbound.packets
ipsec.outbound.security_parameter_index	integer	query	False	Filter by ipsec.outbound.security_parameter_index
ipsec.inbound.packets	integer	query	False	Filter by ipsec.inbound.packets
ipsec.inbound.bytes	integer	query	False	Filter by ipsec.inbound.bytes
ipsec.inbound.security_parameter_index	integer	query	False	Filter by ipsec.inbound.security_parameter_index
ipsec.action	string	query	False	Filter by ipsec.action
name	string	query	False	Filter by name
fields	array[string]	query	False	Specify the fields to return.

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
error	error	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "num_records": "1",
  "records": [
    {
      "cipher_suite": "string",
      "ike": {
        "authentication": "string",
        "state": "string",
        "version": "string"
      },
      "ipsec": {
        "action": "string",
        "state": "string"
      },
      "local_address": "string",
      "name": "string",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
}
```

```

    "remote_address": "string",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "type": "string",
    "uuid": "string"
  }
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

ike

Objects containing parameters specific to IKE (Internet Key Exchange) security association.

Name	Type	Description
authentication	string	Authentication method for internet key exchange protocol.
initiator_security_parameter_index	integer	Initiator's security parameter index for the IKE security association.

Name	Type	Description
responder_security_parameter_index	integer	Responder's security parameter index for the IKE security association.
state	string	State of the IKE connection.
version	string	Internet key exchange protocol version.

inbound

Status for inbound parameters for the IPsec security association.

Name	Type	Description
bytes	integer	Number of inbound bytes for the IPsec security association.
packets	integer	Number of inbound packets for the IPsec security association.
security_parameter_index	integer	Inbound security parameter index for the IPSec security association.

outbound

Status for outbound parameters for the IPsec security association.

Name	Type	Description
bytes	integer	Number of outbound bytes for the IPsec security association.
packets	integer	Number of outbound packets for the IPsec security association.
security_parameter_index	integer	Outbound security parameter index for the IPSec security association.

ipsec

Objects containing parameters specific to IPsec security association.

Name	Type	Description
action	string	Action for the IPsec security association.
inbound	inbound	Status for inbound parameters for the IPsec security association.
outbound	outbound	Status for outbound parameters for the IPsec security association.
state	string	State of the IPsec security association.

[_links](#)

Name	Type	Description
self	href	

node

Node with the security association.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

records

Security association object for IPsec security association and IKE (Internet Key Exchange) security association.

Name	Type	Description
cipher_suite	string	Cipher suite for the security association.
ike	ike	Objects containing parameters specific to IKE (Internet Key Exchange) security association.
ipsec	ipsec	Objects containing parameters specific to IPsec security association.
lifetime	integer	Lifetime for the security association in seconds.
local_address	string	Local address of the security association.
name	string	Policy name for the security association.
node	node	Node with the security association.
remote_address	string	Remote address of the security association.
svm	svm	SVM, applies only to SVM-scoped objects.
type	string	Type of security association, it can be IPsec or IKE (Internet Key Exchange).
uuid	string	Unique identifier of the security association.

Retrieve an IPsec or IKE security association

GET /security/ipsec/security-associations/{uuid}

Introduced In: 9.8

Retrieves a specific IPsec or IKE (Internet Key Exchange) security association.

Related ONTAP commands

- `security ipsec show-ipsecsa`
- `security ipsec show-ikesa`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	UUID of IPsec or IKE security association.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
cipher_suite	string	Cipher suite for the security association.
ike	ike	Objects containing parameters specific to IKE (Internet Key Exchange) security association.
ipsec	ipsec	Objects containing parameters specific to IPsec security association.
lifetime	integer	Lifetime for the security association in seconds.
local_address	string	Local address of the security association.
name	string	Policy name for the security association.
node	node	Node with the security association.
remote_address	string	Remote address of the security association.
svm	svm	

Name	Type	Description
type	string	Type of security association, it can be IPsec or IKE (Internet Key Exchange).
uuid	string	Unique identifier of the security association.

Example response

```
{
  "cipher_suite": "string",
  "ike": {
    "authentication": "string",
    "state": "string",
    "version": "string"
  },
  "ipsec": {
    "action": "string",
    "state": "string"
  },
  "local_address": "string",
  "name": "string",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "remote_address": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "string"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66257118	IPsec SA with the specified UUID was not found.
66257119	IPsec SA with the specified UUID was not found.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

ike

Objects containing parameters specific to IKE (Internet Key Exchange) security association.

Name	Type	Description
authentication	string	Authentication method for internet key exchange protocol.
initiator_security_parameter_index	integer	Initiator's security parameter index for the IKE security association.
responder_security_parameter_index	integer	Responder's security parameter index for the IKE security association.
state	string	State of the IKE connection.
version	string	Internet key exchange protocol version.

inbound

Status for inbound parameters for the IPsec security association.

Name	Type	Description
bytes	integer	Number of inbound bytes for the IPsec security association.
packets	integer	Number of inbound packets for the IPsec security association.
security_parameter_index	integer	Inbound security parameter index for the IPsec security association.

outbound

Status for outbound parameters for the IPsec security association.

Name	Type	Description
bytes	integer	Number of outbound bytes for the IPsec security association.

Name	Type	Description
packets	integer	Number of outbound packets for the IPsec security association.
security_parameter_index	integer	Outbound security parameter index for the IPSec security association.

ipsec

Objects containing parameters specific to IPsec security association.

Name	Type	Description
action	string	Action for the IPsec security association.
inbound	inbound	Status for inbound parameters for the IPsec security association.
outbound	outbound	Status for outbound parameters for the IPsec security association.
state	string	State of the IPsec security association.

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Node with the security association.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage key managers

Security key-managers endpoint overview

Overview

A key manager is a key management solution (software or dedicated hardware) that enables other ONTAP client modules to securely and persistently store keys for various uses. For example, WAFL uses the key management framework to store and retrieve the volume encryption keys that it uses to encrypt/decrypt data on NVE volumes. A key manager can be configured at both cluster scope and SVM, with one key manager allowed per SVM. The key management framework in ONTAP supports two mutually exclusive modes for persisting keys: external and onboard.

When an SVM is configured with external key management, the keys are stored on up to four key servers that are external to the system.

Once external key management is enabled for an SVM, key servers can be added or removed using the `/api/security/key-managers/{uuid}/key-servers` endpoint. See [POST `/security/key-managers/{uuid}/key-servers`] and [DELETE `/security/key-managers/{uuid}/key-servers/{server}`] for more details.

Setting up external key management dictates that the required certificates for securely communicating with the key server are installed prior to configuring the key manager. To install the required client and server_ca certificates, use the `/api/security/certificates/` endpoint.

See [POST `/security/certificates`], [GET `/security/certificates/uuid`] and [DELETE `/security/certificates/{uuid}`] for more details.

When an SVM is configured with the Onboard Key Manager, the keys are stored in ONTAP in wrapped format using a key hierarchy created using the salted hash of the passphrase entered when configuring the Onboard Key Manager. This model fits well for customers who use ONTAP to store their own data.

Examples

Creating an external key manager with 1 key server for a cluster

The example key manager is configured at the cluster-scope with one key server. Note that the UUIDs of the certificates are those that are already installed at the cluster-scope. Note the `return_records=true` query parameter is used to obtain the newly created key manager configuration

```

# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-
managers?return_records=true' -H 'accept: application/hal+json' -d "{
\"external\": { \"client_certificate\": { \"uuid\": \"5fb1701a-d922-11e8-
bfe8-005056bb017d\" }, \"server_ca_certificates\": [ { \"uuid\":
\"827d7d31-d6c8-11e8-b5bf-005056bb017d\" } ], \"servers\": [ { \"server\":
\"10.225.89.33:5696\" } ] } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "815e9462-dc57-11e8-9b2c-005056bb017d",
      "external": {
        "client_certificate": {
          "uuid": "5fb1701a-d922-11e8-bfe8-005056bb017d"
        },
        "server_ca_certificates": [
          {
            "uuid": "827d7d31-d6c8-11e8-b5bf-005056bb017d"
          }
        ],
        "servers": [
          {
            "server": "10.225.89.33:5696"
          }
        ]
      },
      "_links": {
        "self": {
          "href": "/api/security/key-managers/815e9462-dc57-11e8-9b2c-
005056bb017d"
        }
      }
    }
  ]
}

```

Creating an external key manager with 1 primary key server and 2 secondary key servers for a cluster

The example key manager is configured at the cluster-scope with one key server and two secondary key servers. Note that the UUIDs of the certificates are those that are already installed at the cluster-scope. Note the *return_records=true* query parameter is used to obtain the newly created key manager configuration

```
# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-managers?return_records=true' -H 'accept: application/hal+json' -d "{
  \"external\": { \"client_certificate\": { \"uuid\": \"5fb1701a-d922-11e8-bfe8-005056bb017d\" }, \"server_ca_certificates\": [ { \"uuid\": \"827d7d31-d6c8-11e8-b5bf-005056bb017d\" } ], \"servers\": [ { \"server\": \"10.225.89.33:5696\", \"secondary_key_servers\": [ \"1.1.1.1\", \"10.72.204.27:5696\" ] } ] } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "815e9462-dc57-11e8-9b2c-005056bb017d",
      "external": {
        "client_certificate": {
          "uuid": "5fb1701a-d922-11e8-bfe8-005056bb017d"
        },
        "server_ca_certificates": [
          {
            "uuid": "827d7d31-d6c8-11e8-b5bf-005056bb017d"
          }
        ],
        "servers": [
          {
            "server": "10.225.89.33:5696",
            "secondary_key_servers": [
              "1.1.1.1",
              "10.72.204.27:5096"
            ]
          }
        ]
      },
      "_links": {
        "self": {
          "href": "/api/security/key-managers/815e9462-dc57-11e8-9b2c-005056bb017d"
        }
      }
    }
  ]
}
```

```
    }  
  }  
}  
]  
}
```

Creating an external key manager with 1 key server for an SVM

The example key manager is configured at the SVM-scope with one key server. Note that the UUIDs of the certificates are those that are already installed in that SVM. Note the *return_records=true* query parameter is used to obtain the newly created key manager configuration

```
# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-
managers?return_records=true' -H 'accept: application/hal+json' -d "{
  \"svm\": { \"uuid\": \"216e6c26-d6c6-11e8-b5bf-005056bb017d\" },
  \"external\": { \"client_certificate\": { \"uuid\": \"91dcaf7c-dbbd-11e8-
9b2c-005056bb017d\" }, \"server_ca_certificates\": [ { \"uuid\":
\"a4d4b8ba-dbbd-11e8-9b2c-005056bb017d\" } ], \"servers\": [ { \"server\":
\"10.225.89.34:5696\" } ] } }"

# The response:
{
  "num_records": 1,
  "records": [
    {
      "uuid": "80af63f2-dbbf-11e8-9b2c-005056bb017d",
      "svm": {
        "uuid": "216e6c26-d6c6-11e8-b5bf-005056bb017d"
      },
      "external": {
        "client_certificate": {
          "uuid": "91dcaf7c-dbbd-11e8-9b2c-005056bb017d"
        },
        "server_ca_certificates": [
          {
            "uuid": "a4d4b8ba-dbbd-11e8-9b2c-005056bb017d"
          }
        ],
        "servers": [
          {
            "server": "10.225.89.34:5696"
          }
        ]
      },
      "_links": {
        "self": {
          "href": "/api/security/key-managers/80af63f2-dbbf-11e8-9b2c-
005056bb017d"
        }
      }
    }
  ]
}
```

Creating an onboard key manager for a cluster

The following example shows how to create an onboard key manager for a cluster with the onboard key manager configured at the cluster-scope.

```
# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-managers' -H 'accept:
application/hal+json' -d '{ "onboard": { "passphrase": "passphrase" } }'
```

Retrieving the key manager configurations for all clusters and SVMs

The following example shows how to retrieve all configured key managers along with their configurations.

```
# The API:
GET /api/security/key-managers

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers?fields=*' -H
'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "uuid": "2345f09c-d6c9-11e8-b5bf-005056bb017d",
      "scope": "svm",
      "svm": {
        "uuid": "0f22f8f3-d6c6-11e8-b5bf-005056bb017d",
        "name": "vs0"
      },
      "external": {
        "client_certificate": {
          "uuid": "4cb15482-d6c8-11e8-b5bf-005056bb017d",
          "_links": {
            "self": {
              "href": "/api/security/certificates/4cb15482-d6c8-11e8-b5bf-005056bb017d/"
            }
          }
        }
      },
      "server_ca_certificates": [
```

```

    {
      "uuid": "8a17c858-d6c8-11e8-b5bf-005056bb017d",
      "_links": {
        "self": {
          "href": "/api/security/certificates/8a17c858-d6c8-11e8-b5bf-005056bb017d/"
        }
      }
    },
    "servers": [
      {
        "server": "10.2.30.4:5696",
        "timeout": 25,
        "username": "",
        "_links": {
          "self": {
            "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-005056bb017d/key-servers/10.2.30.4:5696/"
          }
        }
      },
      {
        "server": "vs0.local1:3678",
        "timeout": 25,
        "username": "",
        "secondary_key_servers": [
          "1.1.1.1",
          "10.72.204.27:5096"
        ],
        "_links": {
          "self": {
            "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-005056bb017d/key-servers/vs0.local1:3678/"
          }
        }
      }
    ],
    "_links": {
      "self": {
        "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-005056bb017d"
      }
    }
  },

```



```

{
  "uuid": "815e9462-dc57-11e8-9b2c-005056bb017d",
  "scope": "cluster",
  "external": {
    "client_certificate": {
      "uuid": "5fb1701a-d922-11e8-bfe8-005056bb017d",
      "_links": {
        "self": {
          "href": "/api/security/certificates/5fb1701a-d922-11e8-bfe8-005056bb017d/"
        }
      }
    },
    "server_ca_certificates": [
      {
        "uuid": "827d7d31-d6c8-11e8-b5bf-005056bb017d",
        "_links": {
          "self": {
            "href": "/api/security/certificates/827d7d31-d6c8-11e8-b5bf-005056bb017d/"
          }
        }
      }
    ],
    "servers": [
      {
        "server": "10.225.89.33:5696",
        "timeout": 25,
        "username": "",
        "_links": {
          "self": {
            "href": "/api/security/key-managers/815e9462-dc57-11e8-9b2c-005056bb017d/key-servers/10.225.89.33:5696/"
          }
        }
      }
    ],
    "_links": {
      "self": {
        "href": "/api/security/key-managers/815e9462-dc57-11e8-9b2c-005056bb017d"
      }
    }
  }
},

```

```

"num_records": 2,
"_links": {
  "self": {
    "href": "/api/security/key-managers?fields=*"
  }
}
}

```

Retrieving the key manager configurations for all clusters and SVMs (showing Onboard Key Manager)

The following example shows how to retrieve all configured key managers along with their configurations.

```

# The API:
GET /api/security/key-managers

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers?fields=*' -H
'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "uuid": "8ba52e0f-ae22-11e9-b747-005056bb7636",
      "scope": "cluster",
      "onboard": {
        "enabled": true,
        "key_backup": "-----BEGIN
BACKUP-----\n <Backup Data>
\n-----END BACKUP-----\n"
      },
      "volume_encryption": {
        "supported": false,
        "message": "The following nodes do not support volume granular
encryption: ntap-vsim2.",
        "code": 65536935
      },
      "is_default_data_at_rest_encryption_disabled": false
    }
  ],
  "num_records": 1
}

```

Retrieving expensive fields such as, status.code and status.message, associated with a key manager.

These values are not retrieved by default with the 'fields=*' option.

The following example shows how to retrieve the expensive objects associated with a key manager.

```
# The API:
GET /api/security/key-managers

# The call:
curl -X GET "https://<mgmt-ip>/api/security/key-managers?fields=status.message,status.code" -H 'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "uuid": "ac305d46-aef4-11e9-ad3c-005056bb7636",
      "status": {
        "message": "No action needed at this time.",
        "code": 65537200
      },
      "_links": {
        "self": {
          "href": "/api/security/key-managers/ac305d46-aef4-11e9-ad3c-005056bb7636"
        }
      }
    },
    {
      "num_records": 1,
      "_links": {
        "self": {
          "href": "/api/security/key-managers?fields=status.message,status.code"
        }
      }
    }
  ]
}
```

Retrieving a specific key manager configuration

The following example shows how to retrieve a specific key manager configuration.

```
# The API:
GET /api/security/key-managers/{uuid}
```

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/<uuid>?fields=*'
-H 'accept: application/hal+json'

# The response:
{
  "uuid": "2345f09c-d6c9-11e8-b5bf-005056bb017d",
  "scope": "svm",
  "svm": {
    "uuid": "0f22f8f3-d6c6-11e8-b5bf-005056bb017d",
    "name": "vs0"
  },
  "external": {
    "client_certificate": {
      "uuid": "4cb15482-d6c8-11e8-b5bf-005056bb017d",
      "_links": {
        "self": {
          "href": "/api/security/certificates/4cb15482-d6c8-11e8-b5bf-005056bb017d/"
        }
      }
    },
    "server_ca_certificates": [
      {
        "uuid": "8a17c858-d6c8-11e8-b5bf-005056bb017d",
        "_links": {
          "self": {
            "href": "/api/security/certificates/8a17c858-d6c8-11e8-b5bf-005056bb017d/"
          }
        }
      }
    ],
    "servers": [
      {
        "server": "10.2.30.4:5696",
        "timeout": 25,
        "username": "",
        "_links": {
          "self": {
            "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-005056bb017d/key-servers/10.2.30.4:5696/"
          }
        }
      }
    ]
  },
}
```

```
{
  "server": "vs0.local1:3678",
  "timeout": 25,
  "username": "",
  "_links": {
    "self": {
      "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-005056bb017d/key-servers/vs0.local1:3678/"
    }
  }
}
],
"_links": {
  "self": {
    "href": "/api/security/key-managers/2345f09c-d6c9-11e8-b5bf-005056bb017d"
  }
}
}
```

Updating the configuration of an external key manager

The following example shows how to update the `server_ca` configuration of an external key manager.

```
# The API:
PATCH /api/security/key-managers/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>?' -H
'accept: application/hal+json' -d "{ \"external\": {
  \"server_ca_certificates\": [ { \"uuid\": \"23b05c58-d790-11e8-b5bf-005056bb017d\" } ] } }"
```

Updating the passphrase of an Onboard Key Manager

The following example shows how to update the passphrase of a given key manager.

```
# The API:
PATCH /api/security/key-managers/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>?' -H
'accept: application/hal+json' -d "{ \"onboard\": {
\"existing_passphrase\": \"existing_passphrase\", \"passphrase\":
\"new_passphrase\" } }"
```

Synchronizing the passphrase of the Onboard Key Manager on a cluster

The following example shows how to synchronize the passphrase on a cluster where the Onboard Key Manager is already configured.

```
# The API:
PATCH /api/security/key-managers/{uuid}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>?' -H
'accept: application/hal+json' -d "{ \"onboard\": {
\"existing_passphrase\": \"existing_passphrase\",      \"synchronize\": true
} }"
```

Configuring the Onboard Key Manager on a cluster

The following example shows how to configure the Onboard Key Manager on a cluster where the Onboard Key Manager is not configured, but is configured on an MetroCluster partner cluster.

```
# The API:
POST /api/security/key-managers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-
managers?return_records=false' -H 'accept: application/hal+json' -H
"Content-Type: application/json" -d "{ \"onboard\": {      \"passphrase\":
\"passphrase\",      \"synchronize\": true  } }"
```

Deleting a configured key manager

The following example shows how to delete a key manager given its UUID.

```
# The API:
DELETE /api/security/key-managers/{uuid}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/key-managers/<uuid>?' -H
'accept: application/hal+json'
```

Adding a key server to an external key manager

The following example shows how to add a key server with two secondary key servers to an external key manager.

```
# The API:
POST /api/security/key-managers/{uuid}/key-servers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-servers?return_records=true' -H 'accept: application/hal+json' -d '{"server": "10.225.89.34:5696", "secondary_key_servers": ["1.1.1.1", "10.72.204.27:5696"] }'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "server": "10.225.89.34:5696",
      "secondary_key_servers": [
        "1.1.1.1",
        "10.72.204.27:5096"
      ],
      "_links": {
        "self": {
          "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-005056bb017d/key-servers/10.225.89.34%3A5696"
        }
      }
    }
  ]
}
```

Adding 2 key servers to an external key manager

The following example shows how to add 2 key servers to an external key manager. Note that the *records* property is used to add multiple key servers to the key manager in a single API call.


```
# The API:
POST /api/security/key-managers/{uuid}/key-servers

# The call:
curl -X POST 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-servers?return_records=true' -H 'accept: application/hal+json' -d '{"records": [ { "server": "10.225.89.34:5696" }, { "server": "10.225.89.33:5696" } ] }'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-005056bb017d/key-servers/"
        }
      }
    }
  ]
}
```

Retrieving all the key servers configured in an external key manager

The following example shows how to retrieve all key servers configured in an external key manager.

```
# The API:
GET /api/security/key-managers/{uuid}/key-servers

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-servers?fields=*' -H 'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "uuid": "43e0c191-dc5c-11e8-9b2c-005056bb017d",
      "server": "10.225.89.33:5696",
      "timeout": 25,
      "username": "",

```

```

    "secondary_key_servers": [
      "1.1.1.1",
      "10.72.204.27:5096"
    ],
    "_links": {
      "self": {
        "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-005056bb017d/key-servers/10.225.89.33%3A5696"
      }
    }
  },
  {
    "uuid": "43e0c191-dc5c-11e8-9b2c-005056bb017d",
    "server": "10.225.89.34:5696",
    "timeout": 25,
    "username": "",
    "_links": {
      "self": {
        "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-005056bb017d/key-servers/10.225.89.34%3A5696"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-005056bb017d/key-servers?fields=*"
  }
}
}

```

Retrieving a specific key server configured in an external key manager

The following example shows how to retrieve a specific key server configured in an external key manager.

```
# The API:
GET /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X GET 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-servers/{server}?fields=*' -H 'accept: application/hal+json'

# The response:
{
  "uuid": "43e0c191-dc5c-11e8-9b2c-005056bb017d",
  "server": "10.225.89.34:5696",
  "timeout": 25,
  "username": "",
  "_links": {
    "self": {
      "href": "/api/security/key-managers/43e0c191-dc5c-11e8-9b2c-005056bb017d/key-servers/10.225.89.34:5696"
    }
  }
}
```

Updating a specific key server configuration configured in an external key manager

The following example shows how to update a specific key server configured in an external key manager.

```
# The API:
PATCH /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-servers/{server}' -H 'accept: application/hal+json' -d '{" \"timeout\": 45 }'
```

The following example shows how to update the set of secondary key servers associated with a key server.

```
# The API:
PATCH /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-servers/{server}' -H 'accept: application/hal+json' -d "{
  \"secondary_key_servers\": [ \"1.1.1.1\", \"10.72.204.27:5696\" ] }"
```

Deleting a key server from an external key manager

The following example shows how to delete a key server from an external key manager.

```
# The API:
DELETE /api/security/key-managers/{uuid}/key-servers/{server}

# The call:
curl -X DELETE 'https://<mgmt-ip>/api/security/key-managers/<uuid>/key-servers/{server}' -H 'accept: application/hal+json'
```

Retrieve key managers

GET /security/key-managers

Introduced In: 9.6

Retrieves key managers.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `status.message`
- `status.code`

Related ONTAP commands

- `security key-manager show-keystore`
- `security key-manager external show`
- `security key-manager external show-status`
- `security key-manager onboard show-backup`

Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
scope	string	query	False	Filter by scope
volume_encryption.code	integer	query	False	Filter by volume_encryption.code <ul style="list-style-type: none">• Introduced in: 9.7
volume_encryption.supported	boolean	query	False	Filter by volume_encryption.supported <ul style="list-style-type: none">• Introduced in: 9.7
volume_encryption.message	string	query	False	Filter by volume_encryption.message <ul style="list-style-type: none">• Introduced in: 9.7
onboard.key_backup	string	query	False	Filter by onboard.key_backup <ul style="list-style-type: none">• Introduced in: 9.7
onboard.enabled	boolean	query	False	Filter by onboard.enabled
is_default_data_at_rest_encryption_disabled	boolean	query	False	Filter by is_default_data_at_rest_encryption_disabled <ul style="list-style-type: none">• Introduced in: 9.7
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name

Name	Type	In	Required	Description
status.code	integer	query	False	Filter by status.code • Introduced in: 9.7
status.message	string	query	False	Filter by status.message • Introduced in: 9.7
external.server_ca_certificates.name	string	query	False	Filter by external.server_ca_certificates.name • Introduced in: 9.8
external.server_ca_certificates.uuid	string	query	False	Filter by external.server_ca_certificates.uuid
external.client_certificate.name	string	query	False	Filter by external.client_certificate.name • Introduced in: 9.8
external.client_certificate.uuid	string	query	False	Filter by external.client_certificate.uuid
external.servers.server	string	query	False	Filter by external.servers.server
external.servers.username	string	query	False	Filter by external.servers.username
external.servers.secondary_key_servers	string	query	False	Filter by external.servers.secondary_key_servers • Introduced in: 9.8

Name	Type	In	Required	Description
external.servers.timeout	integer	query	False	Filter by external.servers.timeout
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
num_records	integer	Number of records
records	array[security_key_manager]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "external": {
        "client_certificate": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cert1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "server_ca_certificates": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "cert1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          }
        ],
        "servers": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            }
          }
        ]
      }
    }
  ]
}
```



```

"status": {
  "code": "346758",
  "message": "This cluster is part of a MetroCluster
configuration. Use the REST API POST method security/key_managers/ with
the synchronize option and the same passphrase on the partner cluster
before proceeding with any key manager operations. Failure to do so
could lead to switchover or switchback failure."
},
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "string",
"volume_encryption": {
  "code": "346758",
  "message": "No platform support for volume encryption in
following nodes - node1, node2."
}
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

client_certificate

Client certificate

Name	Type	Description
_links	_links	
name	string	Certificate name
uuid	string	Certificate UUID

server_ca_certificates

Name	Type	Description
_links	_links	
name	string	Certificate name
uuid	string	Certificate UUID

self_link

Name	Type	Description
self	href	

key_server_readcreate

Name	Type	Description
_links	self_link	
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
server	string	External key server for key management. If no port is provided, a default port of 5696 is used.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	Username credentials for connecting with the key server.

external

Configures external key management

Name	Type	Description
client_certificate	client_certificate	Client certificate
server_ca_certificates	array[server_ca_certificates]	The UUIDs of the server CA certificates already installed in the cluster or SVM. The array of certificates are common for all the key servers per SVM.
servers	array[key_server_readcreate]	The set of external key servers.

onboard

Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.

Name	Type	Description
enabled	boolean	Is the onboard key manager enabled?
existing_passphrase	string	The cluster-wide passphrase. This is not audited.

Name	Type	Description
key_backup	string	Backup of the onboard key manager's key hierarchy. It is required to save this backup after configuring the onboard key manager to help in the recovery of the cluster in case of catastrophic failures.

status

Optional status information on the current state of the key manager indicating if it is fully setup or requires more action.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns 0 if the setup is complete. This is an advanced property; there is an added cost to retrieving its value. The property is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <code>fields</code> query parameter or GET for all advanced properties is enabled.
message	string	Current state of the key manager indicating any additional steps to perform to finish the setup. This is an advanced property; there is an added cost to retrieving its value. The property is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <code>fields</code> query parameter or GET for all advanced properties is enabled.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

volume_encryption

Indicates whether volume encryption is supported in the cluster.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns a 0 if volume encryption is supported in all nodes of the cluster.
message	string	Reason for not supporting volume encryption.
supported	boolean	Set to true when volume encryption support is available on all nodes of the cluster.

security_key_manager

Name	Type	Description
_links	_links	
external	external	Configures external key management
is_default_data_at_rest_encryption_disabled	boolean	Indicates whether default data-at-rest encryption is disabled in the cluster.
onboard	onboard	Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".

Name	Type	Description
status	status	Optional status information on the current state of the key manager indicating if it is fully setup or requires more action.
svm	svm	
uuid	string	
volume_encryption	volume_encryption	Indicates whether volume encryption is supported in the cluster.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a key manager

POST /security/key-managers

Introduced In: 9.6

Creates a key manager.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create a key manager.
- `external.client_certificate` - Client certificate. Required only when creating an external key manager.

- `external.server_ca_certificates` - Server CA certificates. Required only when creating an external key manager.
- `external.servers.server` - Key servers. Required only when creating an external key manager.
- `onboard.passphrase` - Cluster-wide passphrase. Required only when creating an Onboard Key Manager.
- `synchronize` - Synchronizes missing onboard keys on any node in the cluster. Required only when creating an Onboard Key Manager at the partner site of a MetroCluster configuration.

Related ONTAP commands

- `security key-manager external enable`
- `security key-manager onboard enable`
- `security key-manager onboard sync`

Parameters

Name	Type	In	Required	Description
<code>return_records</code>	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
<code>external</code>	external	Configures external key management
<code>onboard</code>	onboard	Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.
<code>svm</code>	svm	
<code>uuid</code>	string	

Example request

[illegible]

```

AQAAAAAAYAQAAAAAALgePkcAAAAAIgAAAAAAAOAAAAAAAEOTcR0AAAA
AAAAAAAAAACAAAAAAJAGr3tJA/LRzUQRHwv+1aWvAAAAAAAAAACIAAAAAAA
KAAAAAAACIlCHZAAAAAAAAAAAAAAAAAgAAAAAAQCafcabsxRXMM7gxhLRrzh
AAAAAAAAAAkAAAAAAAIAAAAAAAAA2JjQBQAAACt4IqXcNpVggah10axLsN4
yQjnNVKwY7mANB29042hI7b70DTGCTaVAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBLZXkgQmxvYgABAAAAAwAAABgBAAAAAA
7sbaoQAAAAAiAAAAAAACgAAAAAAAQ5NxHQAAAAAAAAAAAAAAAAIAAAAAAka
ave0kD8tHNRBEfC/7Vpa8AAAAAAAAAIgAAAAAAAOAAAAAAALOHfWkAAAAA
AAAAAAAAAACAAAAAABAMoI9UxrHOGthQm/CB+EHdAAAAAAAAAACQAAAAAAAg
AAAAAAAAACnMmUtAAAAAGV8AtPzENFgsGdsFvmucmYrlQCsfew0HDSFKaZqK6
W8IEVzBAhPoAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
-----END BACKUP-----",
    "passphrase": "The cluster password of length 32-256 ASCII
characters.",
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}

```

Response

```

Status: 201, Created

```

Name	Type	Description
num_records	integer	Number of records
records	array[security_key_manager]	

Example response

[illegible]

```

AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAABOZXRBcHAgS2V5IEJsb2IA
AQAAAMAAAYAQAAAAAALgePkcAAAAAIgAAAAAAAAoAAAAAAAEOTcR0AAAA
AAAAAAAAAAACAAAAAAAJAGr3tJA/LRzUQRHwv+1aWvAAAAAAAAAACIAAAAAAAAA
KAAAAAAAAACIlCHZAAAAAAAAAAAAAAAAAAGAAAAAAAAQCafcabsxRXMM7gxhLRrzh
AAAAAAAAAAAKAAAAAAAAAIAAAAAAAAAA2JjQBQAAACt4IqXcNpVggah10axLsN4
yQjnNVKWY7mANB29042hI7b70DTGCTaVAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBLZXkgQmxvYgABAAAAAwAAABgBAAAAAAAA
7sbaoQAAAAAiAAAAAAAAACgAAAAAAAAAQ5NxHQAAAAAAAAAAAAAAAAAIAAAAAKa
ave0kD8tHNRBEfC/7Vpa8AAAAAAAAAAIgAAAAAAAAoAAAAAAALOHfWkAAAAA
AAAAAAAAAAACAAAAAAABAMoI9UxrHOGthQm/CB+EHdAAAAAAAAAACQAAAAAAAA
gAAAAAAAAACnMmUtAAAAAGVk8AtPzENFgsGdsFvnmucmYrlQCsfew0HDSFKaZqK6
W8IEVzBAhPoAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
-----END BACKUP-----",
    "passphrase": "The cluster password of length 32-256 ASCII
characters."
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
]
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
65536038	A maximum of 4 active key servers are allowed.
65536214	Failed to generate cluster key encryption key.
65536216	Failed to add cluster key encryption key.
65536310	Failed to setup the Onboard Key Manager because the MetroCluster peer is unhealthy.
65536508	The platform does not support data at rest encryption.
65536821	The certificate is not installed.

Error Code	Description
65536822	Multitenant key management is not supported in the current cluster version.
65536823	The SVM has key manager already configured.
65536824	Multitenant key management is not supported in MetroCluster configurations.
65536834	Failed to get existing key-server details for the SVM.
65536852	Failed to query supported KMIP protocol versions.
65536870	Key management servers already configured.
65536871	Duplicate key management servers exist.
65536876	External key management requires client and server CA certificates installed and with one or more key servers provided.
65536878	External key management cannot be configured as one or more volume encryption keys of the SVM are stored in cluster key management server.
65536895	External key manager cannot be configured since this cluster is part of a MetroCluster configuration and the partner site of this MetroCluster configuration has Onboard Key Manager configured.
65536900	The Onboard Key Manager cannot be configured because this cluster is part of a MetroCluster configuration and the partner site has the external key manager configured.
65536903	The Onboard Key Manager has failed to configure on some nodes in the cluster. Use the CLI to sync the Onboard Key Manager configuration on failed nodes.
65536906	The Onboard Key Manager has already been configured at the partner site. Use the CLI to sync the Onboard Key Manager with the same passphrase.
65536907	The Onboard Key Manager is already configured. Use the CLI to sync any nodes with the Onboard Key Manager configuration.
65536916	The Onboard Key Manager is only supported for an admin SVM.
65536920	The Onboard Key Manager passphrase length is incorrect.
65537240	The Onboard Key Manager passphrase must be provided when performing a POST/synchronize operation.
65537241	The Onboard Key Manager existing_passphrase must not be provided when performing a POST/synchronize operation.

Error Code	Description
65537244	Unable to sync/create Onboard Key Manager on the local cluster; Onboard Key Manager is already configured on the cluster.
65537245	Unable to sync/create Onboard Key Manager on the local cluster; Onboard Key Manager is not configured on the partner cluster.
65537246	Unable to sync/create Onboard Key Manager on local cluster. This cluster is not part of a MetroCluster configuration.
66060338	Failed to establish secure connection for a key management server due to incorrect server_ca certificates.
66060339	Failed to establish secure connection for a key management server due to incorrect client certificates.
66060340	Failed to establish secure connection for a key management server due to Cryptsoft error.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

client_certificate

Client certificate

Name	Type	Description
name	string	Certificate name
uuid	string	Certificate UUID

server_ca_certificates

Name	Type	Description
name	string	Certificate name
uuid	string	Certificate UUID

self_link

key_server_readcreate

Name	Type	Description
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
server	string	External key server for key management. If no port is provided, a default port of 5696 is used.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	Username credentials for connecting with the key server.

external

Configures external key management

Name	Type	Description
client_certificate	client_certificate	Client certificate
server_ca_certificates	array[server_ca_certificates]	The UUIDs of the server CA certificates already installed in the cluster or SVM. The array of certificates are common for all the key servers per SVM.
servers	array[key_server_readcreate]	The set of external key servers.

onboard

Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.

Name	Type	Description
enabled	boolean	Is the onboard key manager enabled?
existing_passphrase	string	The cluster-wide passphrase. This is not audited.
key_backup	string	Backup of the onboard key manager's key hierarchy. It is required to save this backup after configuring the onboard key manager to help in the recovery of the cluster in case of catastrophic failures.
passphrase	string	The cluster-wide passphrase. This is not audited.

Name	Type	Description
synchronize	boolean	Synchronizes missing onboard keys on any node in the cluster. If a node is added to a cluster that has onboard key management configured, the synchronize operation needs to be performed in a PATCH operation. In a MetroCluster configuration, if onboard key management is enabled on one site, then the synchronize operation needs to be run as a POST operation on the remote site providing the same passphrase.

status

Optional status information on the current state of the key manager indicating if it is fully setup or requires more action.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns 0 if the setup is complete. This is an advanced property; there is an added cost to retrieving its value. The property is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <code>fields</code> query parameter or GET for all advanced properties is enabled.
message	string	Current state of the key manager indicating any additional steps to perform to finish the setup. This is an advanced property; there is an added cost to retrieving its value. The property is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <code>fields</code> query parameter or GET for all advanced properties is enabled.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume_encryption

Indicates whether volume encryption is supported in the cluster.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns a 0 if volume encryption is supported in all nodes of the cluster.
message	string	Reason for not supporting volume encryption.
supported	boolean	Set to true when volume encryption support is available on all nodes of the cluster.

security_key_manager

Name	Type	Description
external	external	Configures external key management
onboard	onboard	Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.
svm	svm	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete key managers

DELETE /security/key-managers/{uuid}

Introduced In: 9.6

Deletes a key manager.

Related ONTAP commands

- security key-manager external disable
- security key-manager onboard disable

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Key manager UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
65536208	Failed to delete the SVM Key ID.

Error Code	Description
65536233	Internal error. Deletion of km_wrapped_kdb key database has failed for the Onboard Key Manager.
65536234	Internal error. Deletion of cluster_kdb key database has failed for the Onboard Key Manager.
65536239	Encrypted volumes are found for the SVM.
65536242	One or more Storage Encryption devices are assigned an authentication key.
65536242	One or more Storage Encryption devices are assigned an authentication key.
65536800	Failed to lookup onboard keys.
65536813	Encrypted kernel core files found.
65536817	Failed to determine if key manager is safe to disable.
65536822	Multitenant key management is not supported in the current cluster version.
65536827	Failed to determine if the SVM has any encrypted volumes.
65536828	External key management is not enabled for the SVM.
65536867	Encrypted volumes are found for the SVM.
196608301	Failed to determine the type of encryption.
196608305	NAE aggregates are found in the cluster.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve key managers

GET /security/key-managers/{uuid}

Introduced In: 9.6

Retrieves key managers.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `status.message`
- `status.code`

Related ONTAP commands

- `security key-manager show-keystore`
- `security key-manager external show`
- `security key-manager external show-status`
- `security key-manager onboard show-backup`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Key manager UUID
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
<code>_links</code>	_links	
<code>external</code>	external	Configures external key management
<code>is_default_data_at_rest_encryption_disabled</code>	boolean	Indicates whether default data-at-rest encryption is disabled in the cluster.

Name	Type	Description
onboard	onboard	Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
status	status	Optional status information on the current state of the key manager indicating if it is fully setup or requires more action.
svm	svm	
uuid	string	
volume_encryption	volume_encryption	Indicates whether volume encryption is supported in the cluster.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "external": {
    "client_certificate": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cert1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "server_ca_certificates": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cert1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    ],
    "servers": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "secondary_key_servers": [
          "string"
        ],
        "server": "keyserver1.com:5698",
        "timeout": "60",
        "username": "admin"
      }
    ]
  },
  "onboard": {
```



```

    "self": {
      "href": "/api/resourcelink"
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string",
  "volume_encryption": {
    "code": "346758",
    "message": "No platform support for volume encryption in following
nodes - node1, node2."
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

client_certificate

Client certificate

Name	Type	Description
_links	_links	
name	string	Certificate name
uuid	string	Certificate UUID

server_ca_certificates

Name	Type	Description
_links	_links	
name	string	Certificate name
uuid	string	Certificate UUID

self_link

Name	Type	Description
self	href	

key_server_readcreate

Name	Type	Description
_links	self_link	
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.

Name	Type	Description
server	string	External key server for key management. If no port is provided, a default port of 5696 is used.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	Username credentials for connecting with the key server.

external

Configures external key management

Name	Type	Description
client_certificate	client_certificate	Client certificate
server_ca_certificates	array[server_ca_certificates]	The UUIDs of the server CA certificates already installed in the cluster or SVM. The array of certificates are common for all the key servers per SVM.
servers	array[key_server_readcreate]	The set of external key servers.

onboard

Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.

Name	Type	Description
enabled	boolean	Is the onboard key manager enabled?
existing_passphrase	string	The cluster-wide passphrase. This is not audited.

Name	Type	Description
key_backup	string	Backup of the onboard key manager's key hierarchy. It is required to save this backup after configuring the onboard key manager to help in the recovery of the cluster in case of catastrophic failures.

status

Optional status information on the current state of the key manager indicating if it is fully setup or requires more action.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns 0 if the setup is complete. This is an advanced property; there is an added cost to retrieving its value. The property is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <code>fields</code> query parameter or GET for all advanced properties is enabled.
message	string	Current state of the key manager indicating any additional steps to perform to finish the setup. This is an advanced property; there is an added cost to retrieving its value. The property is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <code>fields</code> query parameter or GET for all advanced properties is enabled.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

volume_encryption

Indicates whether volume encryption is supported in the cluster.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns a 0 if volume encryption is supported in all nodes of the cluster.
message	string	Reason for not supporting volume encryption.
supported	boolean	Set to true when volume encryption support is available on all nodes of the cluster.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update key managers

PATCH /security/key-managers/{uuid}

Introduced In: 9.6

Updates a key manager.

Required properties

- `onboard.existing_passphrase` - Cluster-wide passphrase. Required only when synchronizing the passphrase of the Onboard Key Manager.
- `synchronize` - Synchronizes missing Onboard Key Manager keys on any node in the cluster. Required only when synchronizing the Onboard Key Manager keys in a local cluster.

Related ONTAP commands

- `security key-manager external modify`
- `security key-manager onboard sync`
- `security key-manager onboard update-passphrase`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Key manager UUID

Request Body

Name	Type	Description
external	external	Configures external key management
is_default_data_at_rest_encryption_disabled	boolean	Indicates whether default data-at-rest encryption is disabled in the cluster.
onboard	onboard	Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.
svm	svm	
uuid	string	

[illegible]

```

KAAAAAAAAACI1CHZAAAAAAAAAAAAAAAAAgAAAAAAQCafcabsxRXMM7gxhLRrzh
AAAAAAAAAAkAAAAAAAAAAIAAAAAAAAAA2JjQBQAAACt4IqXcNpVggah10axLsN4
yQjnNVKWY7mANB29O42hI7b70DTGCTaVAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAE5ldEFwcCBLZXkgQmxvYgABAAAAAwAAABgBAAAAAAA
7sbaoQAAAAAiAAAAAAAAACgAAAAAAAAAQ5NxHQAAAAAAAAAAAAAAAAIAAAAAAka
ave0kD8tHNRBEfC/7Vpa8AAAAAAAAAAIgAAAAAAAAoAAAAAAAAALOHfWkAAAAA
AAAAAAAAAAACAAAAAAABAMoI9UxrHOGthQm/CB+EHdAAAAAAAAAAACQAAAAAAAA
gAAAAAAAAACnMmUtAAAAAGVk8AtPzENFgsGdsFvmucmYrlQCsfew0HDSFKaZqK6
W8IEVzBAhPoAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
-----END BACKUP-----",
  "passphrase": "The cluster password of length 32-256 ASCII
characters."
},
"svm": {
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"uuid": "string"
}

```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
65536139	The existing passphrase value provided does not match the configured passphrase.
65536150	The new passphrase is same as old passphrase.
65536404	The passphrase does not match the accepted length.
65536406	The change of passphrase failed.
65536407	The passphrase update failed on some nodes.
65536408	The passphrase update failed on some nodes.

Error Code	Description
65536802	The passphrase does not match the accepted length in common criteria mode.
65536821	The certificate is not installed.
65536822	Multitenant key management is not supported in the current cluster version.
65536828	External key management is not enabled for the SVM.
65536850	New client certificate public or private keys are different from the existing client certificate.
65536852	Failed to query supported KMIP protocol versions.
65536917	Updating an onboard passphrase requires both new and existing cluster passphrase.
65537242	The Onboard Key Manager existing_passphrase must be provided when performing a PATCH/synchronize operation.
65537243	The Onboard Key Manager passphrase must not be provided when performing a PATCH/synchronize operation.
66060338	Failed to establish secure connection for a key management server due to incorrect server_ca certificates.
66060339	Failed to establish secure connection for a key management server due to incorrect client certificates.
66060340	Failed to establish secure connection for a key management server due to Cryptsoft error.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

client_certificate

Client certificate

Name	Type	Description
name	string	Certificate name
uuid	string	Certificate UUID

server_ca_certificates

Name	Type	Description
name	string	Certificate name
uuid	string	Certificate UUID

self_link

key_server_readcreate

Name	Type	Description
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	Username credentials for connecting with the key server.

external

Configures external key management

Name	Type	Description
client_certificate	client_certificate	Client certificate
server_ca_certificates	array[server_ca_certificates]	The UUIDs of the server CA certificates already installed in the cluster or SVM. The array of certificates are common for all the key servers per SVM.
servers	array[key_server_readcreate]	The set of external key servers.

onboard

Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.

Name	Type	Description
enabled	boolean	Is the onboard key manager enabled?
existing_passphrase	string	The cluster-wide passphrase. This is not audited.
key_backup	string	Backup of the onboard key manager's key hierarchy. It is required to save this backup after configuring the onboard key manager to help in the recovery of the cluster in case of catastrophic failures.
passphrase	string	The cluster-wide passphrase. This is not audited.

Name	Type	Description
synchronize	boolean	Synchronizes missing onboard keys on any node in the cluster. If a node is added to a cluster that has onboard key management configured, the synchronize operation needs to be performed in a PATCH operation. In a MetroCluster configuration, if onboard key management is enabled on one site, then the synchronize operation needs to be run as a POST operation on the remote site providing the same passphrase.

status

Optional status information on the current state of the key manager indicating if it is fully setup or requires more action.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns 0 if the setup is complete. This is an advanced property; there is an added cost to retrieving its value. The property is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <code>fields</code> query parameter or GET for all advanced properties is enabled.
message	string	Current state of the key manager indicating any additional steps to perform to finish the setup. This is an advanced property; there is an added cost to retrieving its value. The property is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <code>fields</code> query parameter or GET for all advanced properties is enabled.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume_encryption

Indicates whether volume encryption is supported in the cluster.

Name	Type	Description
code	integer	Code corresponding to the status message. Returns a 0 if volume encryption is supported in all nodes of the cluster.
message	string	Reason for not supporting volume encryption.
supported	boolean	Set to true when volume encryption support is available on all nodes of the cluster.

security_key_manager

Name	Type	Description
external	external	Configures external key management
is_default_data_at_rest_encryption_disabled	boolean	Indicates whether default data-at-rest encryption is disabled in the cluster.
onboard	onboard	Configures onboard key management. After configuring onboard key management, save the encrypted configuration data in a safe location so that you can use it if you need to perform a manual recovery operation.
svm	svm	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Migrate SVM keys between security key managers

POST /security/key-managers/{source.uuid}/migrate

Introduced In: 9.7

Migrates the keys belonging to an SVM between the cluster's key manager and the SVM's key manager. This operation can run for several minutes.

Required properties

- `source.uuid` - UUID of the source key manager.
- `uuid` - UUID of the destination key manager.

Related ONTAP commands

- `security key-manager migrate`

Parameters

Name	Type	In	Required	Description
source.uuid	string	path	True	Migration source key manager UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
uuid	string	Key manager UUID

Example request

```
{
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563434"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

info

Migration destination key manager UUID

Name	Type	Description
uuid	string	Key manager UUID

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

List key servers configured in an external key manager

GET /security/key-managers/{uuid}/key-servers

Introduced In: 9.6

Retrieves the list of key servers configured in an external key manager.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	External key manager UUID
timeout	integer	query	False	Filter by timeout
secondary_key_servers	string	query	False	Filter by secondary_key_servers <ul style="list-style-type: none">Introduced in: 9.8
server	string	query	False	Filter by server
username	string	query	False	Filter by username
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[key_server]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "password": "password",
      "secondary_key_servers": [
        "string"
      ],
      "server": "keyserver1.com:5698",
      "timeout": "60",
      "username": "username"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

records

Name	Type	Description
_links	_links	
password	string	Password credentials for connecting with the key server. This is not audited.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
server	string	External key server for key management. If no port is provided, a default port of 5696 is used. Not valid in POST if <code>records</code> is provided.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	KMIP username credentials for connecting with the key server.

key_server

Name	Type	Description
_links	_links	
password	string	Password credentials for connecting with the key server. This is not audited.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
server	string	External key server for key management. If no port is provided, a default port of 5696 is used. Not valid in POST if <code>records</code> is provided.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	KMIP username credentials for connecting with the key server.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Add primary key servers to an external key manager

POST /security/key-managers/{uuid}/key-servers

Introduced In: 9.6

Adds key servers to a configured external key manager.

Required properties

- `uuid` - UUID of the external key manager.
- `server` - Key server name.

Related ONTAP commands

- `security key-manager external add-servers`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	External key manager UUID
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
password	string	Password credentials for connecting with the key server. This is not audited.
records	array[records]	An array of key servers specified to add multiple key servers to a key manager in a single API call. Valid in POST only and not valid if <code>server</code> is provided.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.

Name	Type	Description
server	string	External key server for key management. If no port is provided, a default port of 5696 is used. Not valid in POST if records is provided.

Example request

```
{
  "password": "password",
  "records": [
    {
      "password": "password",
      "secondary_key_servers": [
        "string"
      ],
      "server": "keyserver1.com:5698"
    }
  ],
  "secondary_key_servers": [
    "string"
  ],
  "server": "keyserver1.com:5698"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[key_server]	

Example response

```
{
  "records": [
    {
      "password": "password",
      "records": [
        {
          "password": "password",
          "secondary_key_servers": [
            "string"
          ],
          "server": "keyserver1.com:5698"
        }
      ],
      "secondary_key_servers": [
        "string"
      ],
      "server": "keyserver1.com:5698"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
65536038	A maximum of 4 active key servers are allowed.
65536821	The certificate is not installed.
65536822	Multitenant key management is not supported in the current cluster version.
65536824	Multitenant key management is not supported in MetroCluster configurations.
65536828	External key management is not enabled for the SVM.
65536834	Failed to get existing key-server details for the SVM.
65536852	Failed to query supported KMIP protocol versions.
65536870	Key management servers are already configured.

Error Code	Description
65536871	Duplicate key management servers exist.
65536921	The following issues were found. Unable to execute command on KMIP server.
66060338	Unable to establish secure connection to KMIP server due to incorrect server_ca certificates.
66060339	Unable to establish secure connection to KMIP server due to incorrect client certificates.
66060340	Unable to establish secure connection to KMIP server due to Cryptsoft error.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

records

Name	Type	Description
password	string	Password credentials for connecting with the key server. This is not audited.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
server	string	External key server for key management. If no port is provided, a default port of 5696 is used. Not valid in POST if <code>records</code> is provided.

key_server

Name	Type	Description
password	string	Password credentials for connecting with the key server. This is not audited.
records	array[records]	An array of key servers specified to add multiple key servers to a key manager in a single API call. Valid in POST only and not valid if <code>server</code> is provided.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.

Name	Type	Description
server	string	External key server for key management. If no port is provided, a default port of 5696 is used. Not valid in POST if records is provided.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a primary key server

DELETE /security/key-managers/{uuid}/key-servers/{server}

Introduced In: 9.6

Deletes a key server.

Related ONTAP commands

- `security key-manager external remove-servers`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	External key manager UUID

Name	Type	In	Required	Description
server	string	path	True	Key server configured in the external key manager.

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
65536700	The key server contains keys that are currently in use and not available from any other configured key server in the SVM.
65536822	Multitenant key management is not supported in the current cluster version.
65536824	Multitenant key management is not supported in MetroCluster configurations.
65536828	External key management is not enabled for the SVM.
65536843	The key management server is not configured for the SVM.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve key servers configured in an external key manager

GET /security/key-managers/{uuid}/key-servers/{server}

Introduced In: 9.6

Retrieves key servers configured in an external key manager.

Related ONTAP commands

- `security key-manager external show`
- `security key-manager external show-status`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	External key manager UUID
server	string	path	True	Key server configured in the key manager.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
password	string	Password credentials for connecting with the key server. This is not audited.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
server	string	External key server for key management. If no port is provided, a default port of 5696 is used. Not valid in POST if <code>records</code> is provided.
timeout	integer	I/O timeout in seconds for communicating with the key server.

Name	Type	Description
username	string	KMIP username credentials for connecting with the key server.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "password": "password",
  "secondary_key_servers": [
    "string"
  ],
  "server": "keyserver1.com:5698",
  "timeout": "60",
  "username": "username"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

records

Name	Type	Description
_links	_links	
password	string	Password credentials for connecting with the key server. This is not audited.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
server	string	External key server for key management. If no port is provided, a default port of 5696 is used. Not valid in POST if <code>records</code> is provided.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	KMIP username credentials for connecting with the key server.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a primary key server

PATCH /security/key-managers/{uuid}/key-servers/{server}

Introduced In: 9.6

Updates a key server.

Related ONTAP commands

- `security key-manager external modify-server`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	External key manager UUID
server	string	path	True	Key server configured in the external key manager.

Request Body

Name	Type	Description
password	string	Password credentials for connecting with the key server. This is not audited.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.

Name	Type	Description
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	KMIP username credentials for connecting with the key server.

Example request

```
{
  "password": "password",
  "secondary_key_servers": [
    "string"
  ],
  "timeout": "60",
  "username": "username"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
65536822	Multitenant key management is not supported in the current cluster version.
65536824	Multitenant key management is not supported in MetroCluster configurations.
65536828	External key management is not enabled for the SVM.
65536843	The key management server is not configured for the SVM.
65536845	Missing username.
65536846	Missing password.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

records

Name	Type	Description
password	string	Password credentials for connecting with the key server. This is not audited.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	KMIP username credentials for connecting with the key server.

key_server

Name	Type	Description
password	string	Password credentials for connecting with the key server. This is not audited.
secondary_key_servers	array[string]	A list of the secondary key servers associated with the primary key server.
timeout	integer	I/O timeout in seconds for communicating with the key server.
username	string	KMIP username credentials for connecting with the key server.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View and update login message configuration

Security login messages endpoint overview

Overview

You can use this API to display and manage the login messages configuration. The GET request retrieves all of the login messages in the cluster. GET operations on `/security/login/messages/{uuid}` retrieve the login messages configuration by UUID. PATCH operations on `/security/login/messages/{uuid}` update the login messages configuration by UUID.

+

Examples

Retrieving all of the login messages in the cluster

```
# The API:
/api/security/login/messages

# The call:
curl -X GET "https://<mgmt-ip>/api/security/login/messages?fields=*" -H
"accept: application/hal+json"

# The response:
{
```

```

"records": [
  {
    "uuid": "2581e5aa-9fe3-11e8-b309-005056bbef18",
    "scope": "cluster",
    "banner": "*** WARNING: DO NOT PROCEED IF YOU ARE NOT AUTHORIZED!
****\n",
    "message": "#### Welcome to Cluster X ####\n",
    "show_cluster_message": true,
    "_links": {
      "self": {
        "href": "/api/security/login/messages/2581e5aa-9fe3-11e8-b309-
005056bbef18"
      }
    }
  },
  {
    "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
    "scope": "svm",
    "svm": {
      "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
      "name": "svm1"
    },
    "message": "#### Welcome to SVM1 ####\n",
    "show_cluster_message": true,
    "_links": {
      "self": {
        "href": "/api/security/login/messages/7b1b3715-9ffa-11e8-a5dd-
005056bbef18"
      }
    }
  },
  {
    "uuid": "8ddee11e-a58c-11e8-85e0-005056bbef18",
    "scope": "svm",
    "svm": {
      "uuid": "8ddee11e-a58c-11e8-85e0-005056bbef18",
      "name": "svm3"
    },
    "banner": "*** WARNING: This system is for the use of authorized users
only. ****\n",
    "_links": {
      "self": {
        "href": "/api/security/login/messages/8ddee11e-a58c-11e8-85e0-
005056bbef18"
      }
    }
  }
]

```

```
},
{
  "uuid": "f7e41c99-9ffa-11e8-a5dd-005056bbef18",
  "scope": "svm",
  "svm": {
    "uuid": "f7e41c99-9ffa-11e8-a5dd-005056bbef18",
    "name": "svm2"
  },
  "_links": {
    "self": {
      "href": "/api/security/login/messages/f7e41c99-9ffa-11e8-a5dd-005056bbef18"
    }
  }
},
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/security/login/messages?fields=*"
  }
}
}
```

Retrieving the login messages configuration at the cluster scope

```

# The API:
/api/security/login/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/login/messages?scope=cluster&fields=*" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "2581e5aa-9fe3-11e8-b309-005056bbef18",
      "scope": "cluster",
      "banner": "*** WARNING: DO NOT PROCEED IF YOU ARE NOT AUTHORIZED!
****\n",
      "message": "#### Welcome to Cluster X ####\n",
      "show_cluster_message": true,
      "_links": {
        "self": {
          "href": "/api/security/login/messages/2581e5aa-9fe3-11e8-b309-
005056bbef18"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/login/messages?scope=cluster&fields=*"
    }
  }
}

```

Retrieving the login banner configured at the cluster scope

```
# The API:
/api/security/login/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/login/messages?scope=cluster&fields=banner" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "2581e5aa-9fe3-11e8-b309-005056bbef18",
      "scope": "cluster",
      "banner": "*** WARNING: DO NOT PROCEED IF YOU ARE NOT AUTHORIZED!
      ****\n",
      "_links": {
        "self": {
          "href": "/api/security/login/messages/2581e5aa-9fe3-11e8-b309-
005056bbef18"
        }
      }
    },
    ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/login/messages?scope=cluster&fields=banner"
    }
  }
}
```

Retrieving the login messages configuration of a specific SVM

```
# The API:
/api/security/login/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/security/login/messages?svm.name=svm1&fields=*" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
      "scope": "svm",
      "svm": {
        "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
        "name": "svm1"
      },
      "message": "#### Welcome to SVM1 ####\n",
      "show_cluster_message": true,
      "_links": {
        "self": {
          "href": "/api/security/login/messages/7b1b3715-9ffa-11e8-a5dd-
005056bbef18"
        }
      }
    },
    {
      "num_records": 1,
      "_links": {
        "self": {
          "href": "/api/security/login/messages?svm.name=svm1&fields=*"
        }
      }
    }
  ]
}
```

Retrieving the login messages configuration by UUID, including all fields

```
# The API:
/api/security/login/messages/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/security/login/messages/7b1b3715-9ffa-11e8-a5dd-005056bbef18?fields=*" -H "accept: application/hal+json"

# The response:
{
  "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
  "scope": "svm",
  "svm": {
    "uuid": "7b1b3715-9ffa-11e8-a5dd-005056bbef18",
    "name": "svm1"
  },
  "message": "#### Welcome to SVM1 ####\n",
  "show_cluster_message": true,
  "_links": {
    "self": {
      "href": "/api/security/login/messages/7b1b3715-9ffa-11e8-a5dd-005056bbef18"
    }
  }
}
```

Configuring the login banner in a cluster

```
# The API:
/api/security/login/messages

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/login/messages?scope=cluster" -H "accept:
application/hal+json" -H "Content-Type: appplication/json" -d "{
  \"banner\": \"You are entering secure area.\" }"

# The response:
{
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/login/messages?scope=cluster"
    }
  }
}
```

Configuring the message of the day (MOTD) in a cluster

```
# The API:
/api/security/login/messages

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/login/messages?scope=cluster" -H "accept:
application/hal+json" -H "Content-Type: appplication/json" -d "{
  \"message\": \"Welcome to Cluster X\", \"show_cluster_message\": true }"

# The response:
{
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/login/messages?scope=cluster"
    }
  }
}
```

Clearing the login banner and message of the day (MOTD) in a cluster

```
# The API:
/api/security/login/messages

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/login/messages?scope=cluster" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
  \"banner\": \"\", \"message\": \"\" }"

# The response:
{
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/login/messages?scope=cluster"
    }
  }
}
```

Configuring the login messages for a specific SVM

```
# The API:
/api/security/login/messages

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/security/login/messages?svm.name=svm1" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
  \"banner\" : \"AUTHORIZED ACCESS ONLY\", \"message\": \"WELCOME!\" }"

# The response:
{
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/security/login/messages?svm.name=svm1"
    }
  }
}
```

Configuring the login messages by UUID

```
# The API:
/api/security/login/messages/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/login/messages/7b1b3715-
9ffa-11e8-a5dd-005056bbef18" -H "accept: application/hal+json" -H
"Content-Type: application/json" -d "{ \"banner\" : \"AUTHORIZED ACCESS
ONLY\", \"message\": \"WELCOME!\" }"
```

Clearing the login messages configuration by UUID

```
# The API:
/api/security/login/messages/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/login/messages/7b1b3715-9ffa-11e8-a5dd-005056bbef18" -H "accept: application/hal+json" -H "Content-Type: application/json" -d "{ \"banner\": \"\", \"message\": \"\" }"
```

Retrieve login banner and messages of the day

GET /security/login/messages

Introduced In: 9.6

Retrieves the login banner and messages of the day (MOTD) configured in the cluster and in specific SVMs.

Parameters

Name	Type	In	Required	Description
message	string	query	False	Filter by message
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
uuid	string	query	False	Filter by uuid
banner	string	query	False	Filter by banner
scope	string	query	False	Filter by scope
show_cluster_message	boolean	query	False	Filter by show_cluster_message
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[login_messages]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "banner": "string",
      "message": "string",
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "string"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

login_messages

The login banner and message of the day (MOTD) configuration.

Name	Type	Description
_links	_links	

Name	Type	Description
banner	string	<p>The login banner text. This message is displayed during SSH and console device login just before the password prompt displays. When configured, a cluster-level login banner is used for every incoming connection. Each data SVM can override the cluster-level banner to instead display when you log into the SVM. To restore the default setting for a data SVM, set the banner to an empty string.</p> <p>New lines are supplied as either LF or CRLF but are always returned as LF.</p> <p>Optional in the PATCH body.</p>
message	string	<p>The message of the day (MOTD). This message appears just before the clustershell prompt after a successful login. When configured, the cluster message displays first. If you log in as a data SVM administrator, the SVM message is then printed. The cluster-level MOTD can be disabled for a given data SVM using the "show_cluster_message" property.</p> <p>New lines are supplied as either LF or CRLF but are always returned as LF.</p> <p>Optional in the PATCH body.</p>
scope	string	<p>Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".</p>

Name	Type	Description
show_cluster_message	boolean	Specifies whether to show a cluster-level message before the SVM message when logging in as an SVM administrator. This setting can only be modified by the cluster administrator. Optional in the PATCH body.
svm	svm	
uuid	string	The unique identifier (ID) of the login messages configuration.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a login message configuration by UUID

GET /security/login/messages/{uuid}

Introduced In: 9.6

Retrieves the login messages configuration by UUID.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Login messages configuration UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
banner	string	<p>The login banner text. This message is displayed during SSH and console device login just before the password prompt displays. When configured, a cluster-level login banner is used for every incoming connection. Each data SVM can override the cluster-level banner to instead display when you log into the SVM. To restore the default setting for a data SVM, set the banner to an empty string. New lines are supplied as either LF or CRLF but are always returned as LF.</p> <p>Optional in the PATCH body.</p>

Name	Type	Description
message	string	The message of the day (MOTD). This message appears just before the clustershell prompt after a successful login. When configured, the cluster message displays first. If you log in as a data SVM administrator, the SVM message is then printed. The cluster-level MOTD can be disabled for a given data SVM using the "show_cluster_message" property. New lines are supplied as either LF or CRLF but are always returned as LF. Optional in the PATCH body.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
show_cluster_message	boolean	Specifies whether to show a cluster-level message before the SVM message when logging in as an SVM administrator. This setting can only be modified by the cluster administrator. Optional in the PATCH body.
svm	svm	
uuid	string	The unique identifier (ID) of the login messages configuration.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "banner": "string",
  "message": "string",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the login message configuration

PATCH /security/login/messages/{uuid}

Introduced In: 9.6

Updates the login messages configuration.
There are no required fields. An empty body makes no modifications.

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Login messages configuration UUID

Request Body

Name	Type	Description
banner	string	The login banner text. This message is displayed during SSH and console device login just before the password prompt displays. When configured, a cluster-level login banner is used for every incoming connection. Each data SVM can override the cluster-level banner to instead display when you log into the SVM. To restore the default setting for a data SVM, set the banner to an empty string. New lines are supplied as either LF or CRLF but are always returned as LF. Optional in the PATCH body.
message	string	The message of the day (MOTD). This message appears just before the clustershell prompt after a successful login. When configured, the cluster message displays first. If you log in as a data SVM administrator, the SVM message is then printed. The cluster-level MOTD can be disabled for a given data SVM using the "show_cluster_message" property. New lines are supplied as either LF or CRLF but are always returned as LF. Optional in the PATCH body.

Name	Type	Description
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
show_cluster_message	boolean	Specifies whether to show a cluster-level message before the SVM message when logging in as an SVM administrator. This setting can only be modified by the cluster administrator. Optional in the PATCH body.
svm	svm	
uuid	string	The unique identifier (ID) of the login messages configuration.

Example request

```
{
  "banner": "string",
  "message": "string",
  "scope": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "string"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response codes

Error codes	Description
10225636	Only a cluster administrator can modify the <code>show_cluster_message</code> property.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

login_messages

The login banner and message of the day (MOTD) configuration.

Name	Type	Description
banner	string	<p>The login banner text. This message is displayed during SSH and console device login just before the password prompt displays. When configured, a cluster-level login banner is used for every incoming connection. Each data SVM can override the cluster-level banner to instead display when you log into the SVM. To restore the default setting for a data SVM, set the banner to an empty string.</p> <p>New lines are supplied as either LF or CRLF but are always returned as LF.</p> <p>Optional in the PATCH body.</p>

Name	Type	Description
message	string	The message of the day (MOTD). This message appears just before the clustershell prompt after a successful login. When configured, the cluster message displays first. If you log in as a data SVM administrator, the SVM message is then printed. The cluster-level MOTD can be disabled for a given data SVM using the "show_cluster_message" property. New lines are supplied as either LF or CRLF but are always returned as LF. Optional in the PATCH body.
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
show_cluster_message	boolean	Specifies whether to show a cluster-level message before the SVM message when logging in as an SVM administrator. This setting can only be modified by the cluster administrator. Optional in the PATCH body.
svm	svm	
uuid	string	The unique identifier (ID) of the login messages configuration.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage security roles

Security roles endpoint overview

Overview

ONTAP supports Role Based Access Control (RBAC) wherein a user account must be associated with a role and the role defines the privileges and rights for that user account. A privilege defines the access level of the API as either "none", "readonly", or "all". This specifies whether the user account can perform only a GET operation or POST, PATCH, and DELETE operations as well. A role can comprise of multiple tuples and each tuple consists of the REST API and its access level.

For example, "role1" might be a role that has a tuple {"access": "all", "path": "/api/storage/volume"}, which means that a user account with "role1" can perform all GET, POST, PATCH, and DELETE requests on the *api/storage/volume* API or derived APIs which have *api/storage/volume* as the prefix.

In cases where a role has tuples with multiple APIs having the same prefix, the highest match wins out. For example, if "role1" has the following tuples: {"access": "readonly", "path": "/api/cluster"} and {"access": "all", "path": "/api/cluster/schedules"}, then only a GET request is allowed on APIs with *api/cluster* as the prefix; while POST, PATCH and DELETE requests are possible on the *api/cluster/schedules* API.

Predefined (built-in) roles

Related REST APIs are used to form predefined cluster-scoped and SVM-scoped roles, such as: "admin", "backup", "readonly" for cluster and "vsadmin", "vsadmin-backup", "vsadmin-protocol" for SVMs.

These can be retrieved by calling a GET request on */api/security/roles* API and can be assigned to user accounts. See the examples for *api/security/accounts*.

These predefined roles cannot be modified or deleted.

Mapped roles

Before REST APIs, the RBAC roles (legacy roles) were defined to contain the CLI commands and their access levels. Now, almost all REST APIs map to one or more CLI commands. When a role is created using a POST request on */api/security/roles*, a mapped legacy role is created. This legacy role has the same access level (as that of the REST API) for the mapped CLI commands. However, if a legacy role with the same name already exists, the POST operation fails and you need to choose a unique name for the role. The legacy roles cannot be managed using the REST endpoint */api/security/roles* or its derivatives. Legacy roles are managed using the CLI commands "security login role <create | modify | delete> -role <rolename>".

Note that the mapped legacy role (for the REST API role created) cannot be manipulated using the CLI.

The reverse case is not true; the creation of a legacy role will not create a mapped role with equivalent REST APIs.

API restrictions

Numerous APIs are scoped for the cluster level only. This results in an access error if assigned to an SVM-scoped role. For example, */api/cluster/nodes* does not work when added as a tuple entry for an SVM-scoped role.

A number of APIs allowed for an SVM-scoped role might have restrictions on the access level. For example, */api/network/ethernet/ports* cannot have an access level of "all" for an SVM-scoped role; this results in an access error when a POST or PATCH request is made.

Roles created with a REST API path prefix which is common to many APIs might have restrictions based on the scope of the role; cluster or SVM.

For example, {"access":"all","path":"/api/security"} might be a tuple entry for an SVM role. Any GET, POST, PATCH, or DELETE operation fails on API */api/security/accounts* while the same on */api/security/login/messages* succeeds. However, a role with exactly the same tuple when created at the cluster-scope level allows the operations.

Numerous APIs have restrictions on the objects that can be operated on based on the context of the SVM or cluster. For example, a POST request on */api/security/authentication/password* API changes the password for a user account. If executed in the context of an SVM (POST request on an SVM interface), only the password of the user executing the POST can be modified, and attempts to modify the password of any other user results in an access error. However, if a POST request is performed by a cluster administrator account, the password for any user account (cluster or SVM) can be modified.

Examples

Creating a cluster-scoped custom role

Specify the role name and the tuples (of REST APIs and their access level) in the body of the POST request. The owner.uuid or owner.name are not required to be specified for a cluster-scoped role.

```
# The API:
POST "/api/security/roles"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles" -d
'{"name":"cluster_role", "privileges" :
[{"access":"readonly","path":"/api/cluster/jobs"}, {"access":"all","path":"/api/application/applications"}, {"access":"readonly","path":"/api/application/templates"}]}'
```

Creating an SVM-scoped custom role

For an SVM-scoped role, specify either owner.name or owner.uuid in the request body along with other parameters for the role. These correspond to the name or UUID of the SVM for which the role is being created and can be obtained from the response body of the GET request performed on the */api/svm/svms* API.

```
# The API:
POST "/api/security/roles"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles" -d '{"owner": {"uuid": "9f93e553-4b02-11e9-a3f9-005056bb7acd"}, "name": "svm_role", "privileges": [{"access": "readonly", "path": "/api/cluster/jobs"}, {"access": "all", "path": "/api/application/applications"}, {"access": "readonly", "path": "/api/application/templates"}]}'
```

Retrieving the configured roles

All of the roles or a filtered list of roles (for example by name, predefined, and so on) can be retrieved.

```
# The API:
GET "/api/security/roles"

# The call to retrieve all the roles configured in the cluster:
curl -X GET "https://<mgmt-ip>/api/security/roles"

# The response:
{
  "records": [
    {
      "owner": {
        "uuid": "2903de6f-4bd2-11e9-b238-0050568e2e25",
        "name": "cluster1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/2903de6f-4bd2-11e9-b238-0050568e2e25"
          }
        }
      },
      "name": "admin",
      "privileges": [
        {
          "path": "/api",
          "access": "all",
          "_links": {
            "self": {
              "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-0050568e2e25/admin/privileges/%2Fapi"
            }
          }
        }
      ]
    }
  ]
}
```



```

    }
  ],
  "builtin": true,
  "scope": "cluster",
  "_links": {
    "self": {
      "href": "/api/security/roles/2903de6f-4bd2-11e9-b238-0050568e2e25/admin"
    }
  }
},
{
  "owner": {
    "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
      }
    }
  },
  "name": "vsadmin",
  "privileges": [
    {
      "path": "/api/application/applications",
      "access": "all",
      "_links": {
        "self": {
          "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/vsadmin/privileges/%2Fapi%2Fapplication%2Fapplications"
        }
      }
    },
    {
      "path": "/api/application/templates",
      "access": "readonly",
      "_links": {
        "self": {
          "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/vsadmin/privileges/%2Fapi%2Fapplication%2Ftemplates"
        }
      }
    }
  ],
  {
    "path": "/api/cluster",
    "access": "readonly",

```

```

    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/vsadmin/privileges/%2Fapi%2Fcluster"
      }
    },
    {
      "path": "/api/svm/svms",
      "access": "readonly",
      "_links": {
        "self": {
          "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/vsadmin/privileges/%2Fapi%2Fsvm%2Fsvms"
        }
      }
    },
    {
      "path": "/api/svms",
      "access": "readonly",
      "_links": {
        "self": {
          "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/vsadmin/privileges/%2Fapi%2Fsvms"
        }
      }
    },
    "builtin": true,
    "scope": "svm",
    "_links": {
      "self": {
        "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/vsadmin"
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/security/roles"
    }
  }
}

```

Using a scoped call to retrieve the configured roles

```
# Scoped call to retrieve all the roles for a particular SVM using
owner.uuid:
curl -X GET "https://<mgmt-ip>/api/security/roles/?owner.uuid=aaef7c38-
4bd3-11e9-b238-0050568e2e25"

# Scoped call to retrieve all the roles for a particular SVM using
owner.name:
curl -X GET "https://<mgmt-ip>/api/security/roles/?owner.name=svm1"

# Scoped call to retrieve the roles having vsadmin as the prefix in the
role name:
curl -X GET "https://<mgmt-ip>/api/security/roles/?name=vsadmin*"

# Scoped call to retrieve the predefined roles:
curl -X GET "https://<mgmt-ip>/api/security/roles/?builtin=true"

# Scoped call to retrieve the custom roles:
curl -X GET "https://<mgmt-ip>/api/security/roles/?builtin=false"
```

Retrieve a list of roles configured in the cluster

GET /security/roles

Introduced In: 9.6

Retrieves a list of roles configured in the cluster.

Related ONTAP commands

- `security login rest-role show`

Learn more

- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
privileges.path	string	query	False	Filter by privileges.path <ul style="list-style-type: none">• Introduced in: 9.7

Name	Type	In	Required	Description
privileges.access	string	query	False	Filter by privileges.access • Introduced in: 9.7
scope	string	query	False	Filter by scope • Introduced in: 9.7
owner.uuid	string	query	False	Filter by owner.uuid • Introduced in: 9.7
owner.name	string	query	False	Filter by owner.name • Introduced in: 9.7
name	string	query	False	Filter by name • Introduced in: 9.7
builtin	boolean	query	False	Filter by builtin • Introduced in: 9.7
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[role]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "admin",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "privileges": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "access": "readonly",
          "path": "/api/storage/volumes"
        }
      ],
      "scope": "string"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

owner

Owner name and UUID that uniquely identifies the role.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
_links	_links	
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

role

A named set of privileges that defines the rights an account has when it is assigned the role.

Name	Type	Description
_links	_links	
builtin	boolean	Indicates if this is a built-in (pre-defined) role which cannot be modified or deleted.
name	string	Role name
owner	owner	Owner name and UUID that uniquely identifies the role.
privileges	array[role_privilege]	The list of privileges that this role has been granted.
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a new cluster-scoped or SVM-scoped role

POST /security/roles

Introduced In: 9.6

Creates a new cluster-scoped role or an SVM-scoped role. For an SVM-scoped role, specify either the SVM name as the `owner.name` or SVM UUID as the `owner.uuid` in the request body along with other parameters for the role. The `owner.uuid` or `owner.name` are not required to be specified for a cluster-scoped role.

Required parameters

- `name` - Name of the role to be created.
- `privileges` - Array of privilege tuples. Each tuple consists of a REST API path and its desired access level.

Optional parameters

- `owner.name` or `owner.uuid` - Name or UUID of the SVM for an SVM-scoped role.

Related ONTAP commands

- `security login rest-role create`

Learn more

- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
builtin	boolean	Indicates if this is a built-in (pre-defined) role which cannot be modified or deleted.
name	string	Role name
owner	owner	Owner name and UUID that uniquely identifies the role.
privileges	array[role_privilege]	The list of privileges that this role has been granted.

Name	Type	Description
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

Example request

```
{
  "name": "admin",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "privileges": [
    {
      "access": "readonly",
      "path": "/api/storage/volumes"
    }
  ],
  "scope": "string"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2621462	The supplied SVM does not exist.
5636129	Role with given name has not been defined.
5636143	Vserver admin cannot use the API with this access level.
5636144	Invalid value specified for access level.
5636169	Invalid character in URI.
5636170	URI does not exist.

Error Code	Description
5636171	Role already exists in legacy role table.
13434890	Vserver-ID failed for Vserver roles.
13434891	UUID lookup failed for Vserver roles.
13434892	Roles is a required field.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

owner

Owner name and UUID that uniquely identifies the role.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

role

A named set of privileges that defines the rights an account has when it is assigned the role.

Name	Type	Description
builtin	boolean	Indicates if this is a built-in (pre-defined) role which cannot be modified or deleted.
name	string	Role name
owner	owner	Owner name and UUID that uniquely identifies the role.
privileges	array[role_privilege]	The list of privileges that this role has been granted.

Name	Type	Description
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View or delete a role

Security roles owner.uuid name endpoint overview

Overview

This API is used to retrieve or delete a role. The role can be SVM-scoped or cluster-scoped.

Specify the owner UUID and the role name in the URI path. The owner UUID corresponds to the UUID of the SVM for which the role has been created and can be obtained from the response body of a GET call performed on one of the following APIs:

/api/security/roles for all roles

/api/security/roles/?scope=svm for SVM-scoped roles

/api/security/roles/?owner.name={svm-name} for roles in a specific SVM

This API response contains the complete URI for each role that can be used for retrieving or deleting a role.



The pre-defined roles can be retrieved but cannot be deleted.

Examples

Retrieving a role configuration

```
# The API:
GET "/api/security/roles/{owner.uuid}/{name}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/secure_role"

# The response:
{
  "owner": {
    "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/aaef7c38-4bd3-11e9-b238-0050568e2e25"
      }
    }
  },
  "name": "secure_role",
  "privileges": [
    {
      "path": "/api/security",
      "access": "all",
      "_links": {
        "self": {
          "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/secure_role/privileges/%2Fapi%2Fsecurity"
        }
      }
    }
  ],
  "builtin": false,
  "scope": "svm",
  "_links": {
    "self": {
      "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/secure_role"
    }
  }
}
```

Deleting a custom role

```
# The API:
DELETE "/api/security/roles/{owner.uuid}/{name}"

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1"
```

Delete a role

DELETE /security/roles/{owner.uuid}/{name}

Introduced In: 9.6

Deletes the specified role.

Required parameters

- `name` - Name of the role to be deleted.
- `owner.uuid` - UUID of the SVM housing the role.

Related ONTAP commands

- `security login rest-role delete`

Learn more

- [DOC /security/roles/{owner.uuid}/{name}](#)
- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Role owner UUID
name	string	path	True	Role name to be deleted.

Response

```
Status: 200, Ok
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5636172	User accounts detected with this role assigned. Update or delete those accounts before deleting this role.
5636173	Features require an effective cluster version of 9.6 or later.
13434890	Vserver-ID failed for Vserver roles.
13434893	The SVM does not exist.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the details of a role

GET /security/roles/{owner.uuid}/{name}

Introduced In: 9.6

Retrieves the details of the specified role.

Related ONTAP commands

- `security login rest-role show`

Learn more

- [DOC /security/roles/{owner.uuid}/{name}](#)
- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Role owner UUID
name	string	path	True	Role name

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
builtin	boolean	Indicates if this is a built-in (pre-defined) role which cannot be modified or deleted.
name	string	Role name
owner	owner	Owner name and UUID that uniquely identifies the role.
privileges	array[role_privilege]	The list of privileges that this role has been granted.
scope	string	Scope of the entity. Set to "cluster" for cluster owned objects and to "svm" for SVM owned objects.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "admin",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "privileges": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access": "readonly",
      "path": "/api/storage/volumes"
    }
  ],
  "scope": "string"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

owner

Owner name and UUID that uniquely identifies the role.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

role_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
_links	_links	
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage role privilege details

Security roles owner.uuid name privileges endpoint overview

Overview

This API is used to configure the role privileges (tuples of REST URI path and its access levels). It also retrieves all of the privilege tuples for a role and can add a tuple to an existing role.

The role can be SVM-scoped or cluster-scoped.

Specify the owner UUID and the role name in the URI path. The owner UUID corresponds to the UUID of the SVM for which the role has been created and can be obtained from the response body of a GET request performed on one of the following APIs:
`/api/security/roles` for all the roles
`/api/security/roles/?scope=svm` for SVM-scoped roles
`/api/security/roles/?owner.name=<svm-name><i></i>` for roles in a specific SVM
This API response contains the complete URI for each role and can be used after suffixing it with `_<i></i>_` `_<i></i>_`



The pre-defined roles can be retrieved but cannot be updated.

Examples

Adding a privilege tuple for an existing custom role

```
# The API:
POST "/security/roles/{owner.uuid}/{name}/privileges"

# The call:
curl -X POST "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1/privileges" -d
'{"access":"readonly","path":"/api/protocols"}'
```

Retrieving all the privilege tuples for a role

```
# The API:
GET "/api/security/roles/{owner.uuid}/{name}/privileges"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1/privileges"

# The response:
{
  "records": [
    {
      "path": "/api/application",
      "access": "all",
      "_links": {
        "self": {
          "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fapplication"
        }
      }
    },
    {
      "path": "/api/protocols",
      "access": "readonly",
      "_links": {
        "self": {
          "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1/privileges"
    }
  }
}
```

Retrieve privilege details of the specified role

GET /security/roles/{owner.uuid}/{name}/privileges

Introduced In: 9.6

Retrieves privilege details of the specified role.

Related ONTAP commands

- `security login rest-role show`

Learn more

- [DOC /security/roles/{owner.uuid}/{name}/privileges](#)
- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Role owner UUID
name	string	path	True	Role name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[role_privilege]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "access": "readonly",
      "path": "/api/storage/volumes"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

role_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
_links	_links	
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Add a privilege tuple to an existing role

`POST /security/roles/{owner.uuid}/{name}/privileges`

Introduced In: 9.6

Adds a privilege tuple (of REST URI and its access level) to an existing role.

Required parameters

- `owner.uuid` - UUID of the SVM that houses this role.
- `name` - Name of the role to be updated.
- `path` - REST URI path (example: `"/api/storage/volumes"`).
- `access` - Desired access level for the REST URI path (one of "all", "readonly" or "none").

Related ONTAP commands

- `security login rest-role create`

Learn more

- [DOC /security/roles/{owner.uuid}/{name}/privileges](#)
- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Role owner UUID
name	string	path	True	Role name
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

Example request

```
{
  "access": "readonly",
  "path": "/api/storage/volumes"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5636129	A role with given name has not been defined.
5636143	A Vserver admin cannot use the API with this access level.
5636144	The value specified for the access level is not valid.
5636169	A character in the URI is not valid.
5636170	The URI does not exist.
5636173	This feature requires an effective cluster version of 9.6 or later.
5636175	Vserver admin cannot have access to given API.
13434890	Vserver-ID failed for Vserver roles.
13434891	UUID LookUp failed for Vserver roles.
13434892	Roles is a required field.

Error Code	Description
13434893	The SVM does not exist.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

role_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage role privilege path

Security roles owner.uuid name privileges path endpoint overview

Overview

A role can comprise of multiple tuples and each tuple consists of the REST API path and its access level. These APIs can be used to retrieve and modify the access level or delete one of the constituent REST API paths within a role.

The role can be SVM-scoped or cluster-scoped.

Specify the owner UUID and the role name in the URI path. The owner UUID corresponds to the UUID of the SVM for which the role has been created and can be obtained from the response body of a GET request performed on one of the following APIs:
 /api/security/roles for all roles
 /api/security/roles/?scope=svm for SVM-scoped roles
 /api/security/roles/?owner.name=<svm-name><i></i> for roles in a specific SVM This API response contains the complete URI for each tuple of the role and can be used for GET, PATCH, or DELETE operations.</svm-name>



The access level for paths in pre-defined roles cannot be updated.

Examples

Updating the access level for a path in the privilege tuple of an existing role

```
# The API:
PATCH "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-
b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols" -d
'{"access":"all"}'
```

Retrieving the access level for a path in the privilege tuple of an existing role


```
# The API:
GET "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols"

# The response:
{
  "owner": {
    "uuid": "aaef7c38-4bd3-11e9-b238-0050568e2e25"
  },
  "name": "svm_role1",
  "path": "/api/protocols",
  "access": "all",
  "_links": {
    "self": {
      "href": "/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols"
    }
  }
}
```

Deleting a privilege tuple from an existing role

```
# The API:
DELETE "/api/security/roles/{owner.uuid}/{name}/privileges/{path}"

# The call:
curl -X DELETE "https://<mgmt-ip>/api/security/roles/aaef7c38-4bd3-11e9-b238-0050568e2e25/svm_role1/privileges/%2Fapi%2Fprotocols"
```

Delete a privilege tuple from the role

```
DELETE /security/roles/{owner.uuid}/{name}/privileges/{path}
```

Introduced In: 9.6

Deletes a privilege tuple (of REST URI and its access level) from the role.

Required parameters

- `owner.uuid` - UUID of the SVM which houses this role.
- `name` - Name of the role to be updated.
- `path` - Constituent REST API path to be deleted from this role.

Related ONTAP commands

- security login rest-role delete

Learn more

- [DOC /security/roles/{owner.uuid}/{name}/privileges/{path}](#)
- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Role owner UUID
name	string	path	True	Role name
path	string	path	True	REST API path

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5636172	User accounts detected with this role assigned. Update or delete those accounts before deleting this role.
5636173	This feature requires an effective cluster version of 9.6 or later.
13434890	Vserver-ID failed for Vserver roles.
13434893	The SVM does not exist.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the access level for a REST API path or command/command directory path for a role

GET /security/roles/{owner.uuid}/{name}/privileges/{path}

Introduced In: 9.6

Retrieves the privilege level for a REST API path for the specified role.

Related ONTAP commands

- `security login rest-role show`

Learn more

- [DOC /security/roles/{owner.uuid}/{name}/privileges/{path}](#)
- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Role owner UUID
name	string	path	True	Role name
path	string	path	True	REST API path
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
access	string	Access level for the REST endpoint.
path	string	REST URI/endpoint

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "access": "readonly",
  "path": "/api/storage/volumes"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the access level for a REST API path or command/command directory path

PATCH /security/roles/{owner.uuid}/{name}/privileges/{path}

Introduced In: 9.6

Updates the privilege level for a REST API path.

Required parameters

- `owner.uuid` - UUID of the SVM that houses this role.
- `name` - Name of the role to be updated.
- `path` - Constituent REST API path whose access level has to be updated.

- `access` - Access level for the path (one of "all", "readonly", or "none")

Related ONTAP commands

- `security login rest-role modify`

Learn more

- [DOC /security/roles/{owner.uuid}/{name}/privileges/{path}](#)
- [DOC /security/roles](#)

Parameters

Name	Type	In	Required	Description
owner.uuid	string	path	True	Role owner UUID
name	string	path	True	Role name
path	string	path	True	REST API path

Request Body

Name	Type	Description
access	string	Access level for the REST endpoint.

Example request

```
{
  "access": "readonly"
}
```

Response

Status: 200, Ok

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

role_privilege

A tuple containing the REST endpoint and the access level assigned to that endpoint.

Name	Type	Description
access	string	Access level for the REST endpoint.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage SSH server

Security SSH endpoint overview

Overview

ONTAP supports SSH server that can be accessed from any standard SSH client. A user account needs to be

associated with SSH as the application (refer the documentation for [api/security/accounts DOC](#) [/security/accounts](#) . Upon connecting from a client, the user is authenticated and a command line shell is presented.

This endpoint is used to retrieve or modify the SSH configuration at the cluster level. The configuration consists of SSH security parameters (security algorithms and maximum authentication retry attempts allowed before closing the connection) and SSH connection limits.

The security algorithms include SSH key exchange algorithms, ciphers for payload encryption, and MAC algorithms. This configuration is the default for all newly created SVMs; existing SVM configurations are not impacted.

The SSH connection limits include maximum connections per second, maximum simultaneous sessions from the same client host, and overall maximum SSH connections at any given point in time. The connection limits are per node and will be the same for all nodes in the cluster.

Examples

Updating the SSH security parameters

Specify the algorithms in the body of the PATCH request.

```
# The API:
PATCH "/api/security/ssh"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/ssh" -d '{ "ciphers": [
"aes256_ctr", "aes192_ctr" ], "key_exchange_algorithms": [
"diffie_hellman_group_exchange_sha256", "diffie_hellman_group14_sha1" ],
"mac_algorithms": [ "hmac_sha2_512_etm", "umac_128_etm" ],
"max_authentication_retry_count": 3 }'
```

Updating the SSH connection limits

Specify the connection limits in the body of the PATCH request.

```
# The API:
PATCH "/api/security/ssh"

# The call:
curl -X PATCH "https://<mgmt-ip>/api/security/ssh" -d '{
"connections_per_second": 8, "max_instances": 10, "per_source_limit": 5 }'
```

Retrieving the cluster SSH server configuration

```

# The API:
GET "/api/security/ssh"

# The call:
curl -X GET "https://<mgmt-ip>/api/security/ssh"

# The response:
{
  "ciphers": [
    "aes256_ctr",
    "aes192_ctr"
  ],
  "key_exchange_algorithms": [
    "diffie_hellman_group_exchange_sha256",
    "diffie_hellman_group14_sha1"
  ],
  "mac_algorithms": [
    "hmac_sha2_512_etm",
    "umac_128_etm"
  ],
  "max_authentication_retry_count": 3,
  "connections_per_second": 8,
  "max_instances": 10,
  "per_source_limit": 5,
  "_links": {
    "self": {
      "href": "/api/security/ssh"
    }
  }
}

```

Retrieve cluster SSH server ciphers, MAC algorithms, key exchange algorithms, and connection limits

GET /security/ssh

Introduced In: 9.7

Retrieves the cluster SSH server ciphers, MAC algorithms, key exchange algorithms, and connection limits.

Related ONTAP commands

- security ssh
- security protocol ssh

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
ciphers	array[string]	Ciphers for encrypting the data.
connections_per_second	integer	Maximum connections allowed per second.
key_exchange_algorithms	array[string]	Key exchange algorithms.
mac_algorithms	array[string]	MAC algorithms.
max_authentication_retry_count	integer	Maximum authentication retries allowed before closing the connection.
max_instances	integer	Maximum possible simultaneous connections.
per_source_limit	integer	Maximum connections from the same client host.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "ciphers": [
    "aes256_ctr",
    "aes192_ctr",
    "aes128_ctr"
  ],
  "key_exchange_algorithms": [
    "diffie_hellman_group_exchange_sha256",
    "diffie_hellman_group14_sha1"
  ],
  "mac_algorithms": [
    "hmac_sha1",
    "hmac_sha2_512_etm"
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the SSH server setting for a cluster

PATCH /security/ssh

Introduced In: 9.7

Updates the SSH server setting for a cluster.

Optional parameters

- `ciphers` - Encryption algorithms for the payload
- `key_exchange_algorithms` - SSH key exchange algorithms
- `mac_algorithms` - MAC algorithms

- `max_authentication_retry_count` - Maximum authentication retries allowed before closing the connection
- `connections_per_second` - Maximum allowed connections per second
- `max_instances` - Maximum allowed connections per node
- `per_source_limit` - Maximum allowed connections from the same client host

Related ONTAP commands

- `security ssh`
- `security protocol ssh`

Parameters

Name	Type	In	Required	Description
<code>return_records</code>	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
<code>ciphers</code>	array[string]	Ciphers for encrypting the data.
<code>connections_per_second</code>	integer	Maximum connections allowed per second.
<code>key_exchange_algorithms</code>	array[string]	Key exchange algorithms.
<code>mac_algorithms</code>	array[string]	MAC algorithms.
<code>max_authentication_retry_count</code>	integer	Maximum authentication retries allowed before closing the connection.
<code>max_instances</code>	integer	Maximum possible simultaneous connections.
<code>per_source_limit</code>	integer	Maximum connections from the same client host.

Example request

```
{
  "ciphers": [
    "aes256_ctr",
    "aes192_ctr",
    "aes128_ctr"
  ],
  "key_exchange_algorithms": [
    "diffie_hellman_group_exchange_sha256",
    "diffie_hellman_group14_sha1"
  ],
  "mac_algorithms": [
    "hmac_sha1",
    "hmac_sha2_512_etm"
  ]
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
10682372	There must be at least one key exchange algorithm associated with the SSH configuration.
10682373	There must be at least one cipher associated with the SSH configuration.
10682375	Failed to modify SSH key exchange algorithms.
10682378	Failed to modify SSH ciphers.
10682399	Key exchange algorithm not supported in FIPS enabled mode.
10682400	Failed to modify SSH MAC algorithms.
10682401	MAC algorithm not supported in FIPS enabled mode.

Error Code	Description
10682403	There must be at least one MAC algorithm with the SSH configuration.
10682413	Failed to modify maximum authentication retry attempts.
10682413	Failed to modify maximum authentication retry attempts.
10682418	Cipher not supported in FIPS enabled mode.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

cluster_ssh_server

Name	Type	Description
ciphers	array[string]	Ciphers for encrypting the data.
connections_per_second	integer	Maximum connections allowed per second.
key_exchange_algorithms	array[string]	Key exchange algorithms.
mac_algorithms	array[string]	MAC algorithms.
max_authentication_retry_count	integer	Maximum authentication retries allowed before closing the connection.
max_instances	integer	Maximum possible simultaneous connections.
per_source_limit	integer	Maximum connections from the same client host.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

SnapLock

SnapLock overview

SnapLock is an alternative to the traditional optical "write once, read many" (WORM) data. SnapLock is used for the storage of read-only WORM data. SnapLock is a license-based, disk-based, open-protocol feature that works with application software to administer non-rewritable storage of data. The primary objective of this feature is to provide storage-enforced WORM and retention functionality. SnapLock can be deployed for protecting data in strict regulatory environments in such a way that even the storage administrator is considered an untrusted party. SnapLock provides special purpose volumes in which files can be stored and committed to a non-erasable, non-rewritable state either forever or for a designated retention period. SnapLock allows this retention to be performed at the granularity of individual files. These APIs allow you to manage the following endpoints:

- SnapLock log
- SnapLock compliance-clock
- SnapLock event-retention
- SnapLock file retention
- SnapLock privileged-delete
- SnapLock file fingerprint
- SnapLock legal-hold

APIs

SnapLock log

These APIs are used to create and initialize the SnapLock log configuration for an SVM, modify attributes associated with the SnapLock log configuration, and delete the auditlog configuration.

SnapLock compliance-clock

This API manages the ComplianceClock of the system. ComplianceClock determines the expiry time of the SnapLock objects in the system. ComplianceClock can be initialized only once by the user and once it is set, it cannot be changed or altered by the user.

SnapLock event-retention

This API is used to create Event Based Retention (EBR) policies for an SVM. A policy consists of a policy-name and a retention-period. Only a user with security login role vsadmin-snaplock is allowed to perform this operation.

SnapLock file retention

This API is used to set the retention time of a file. Retention time expects the date in ISO 8601 format or "infinite".

SnapLock privileged-delete

This API is used to perform a privileged-delete operation on unexpired WORM files on a SnapLock Enterprise volume. The only built-in role that has access to the command is "vsadmin-snaplock".

SnapLock file fingerprint

This API is used to start the fingerprint computation on a file.

SnapLock legal-hold

This API is used to retain Compliance-mode WORM files for the duration of a litigation. A file under a Legal-Hold behaves like a WORM file with an indefinite retention period. Only a user with security login role vsadmin-snaplock is allowed to perform the operation.

Manage SnapLock audit logs

Storage SnapLock audit-logs endpoint overview

The SnapLock log volume can be a SnapLock Compliance volume or SnapLock Enterprise volume. The SnapLock log infrastructure creates directories and files in this volume to store the SnapLock log records. The maximum log size specifies the maximum size of a log file that stores SnapLock log records. When the file reaches this size, the log infrastructure archives it and creates a new log file. The default retention period is the length of time the log file is retained, if the SnapLock log records that are stored in the file do not carry any retention period.

Examples

1. Verifies that the audit log is configured for the specified SVM:

```
GET "/api/storage/snaplock/audit-logs/?svm.name=VS0"
```

1. Verifies that the specified volume is an audit log volume:

```
GET "/api/storage/snaplock/audit-logs/?log_volume.volume.name=VS0_ALOG"
```

Examples

1. Creates a SnapLock log configuration by providing SVM name:

```
POST "/api/storage/snaplock/audit-logs" '{"svm": {"name": "VS3"},  
"log_volume": { "volume": { "name": "VS3_ALOG"}, "max_log_size": "20971520",  
"retention_period": "P30Y" } }'
```

1. Creates a SnapLock log configuration by providing SVM UUID:

```
POST "/api/storage/snaplock/audit-logs" '{"svm": {"uuid": "bc744cc7-296d-11e9-a26f-0050568e5b05"}, "log_volume": { "volume": { "name": "VS3_ALOG"}, "max_log_size": "20971520", "retention_period": "P30Y" }}'
```

1. Creates a SnapLock log configuration without specifying a retention period:

```
POST "/api/storage/snaplock/audit-logs" '{"svm": {"name": "VS3"}, "log_volume": {"volume": {"name": "VS3_ALOG"}}}'
```

Examples

1. Updates the audit log volume:

```
PATCH "/api/storage/snaplock/audit-logs/bc744cc7-296d-11e9-a26f-0050568e5b05" '{"log_volume": {"volume": {"name": "VS4_ALOG_NEW"}}}'
```

1. Updates the maximum size of the log file and the retention period:

```
PATCH "/api/storage/snaplock/audit-logs/420cac7a-296a-11e9-a26f-0050568e5b05" '{"log_volume": {"max_log_size": "20971520", "retention_period": "P1Y"}}'
```

1. Archives all of the audit log files:

```
PATCH "/api/storage/snaplock/audit-logs/c7e4fa7d-2968-11e9-a26f-0050568e5b05" '{"log_archive": {"archive": "true"}}'
```

1. Archives the specified audit log file:

```
PATCH "/api/storage/snaplock/audit-logs/c7e4fa7d-2968-11e9-a26f-0050568e5b05" '{"log_archive": {"archive": "true", "base_name": "privileged_delete"}}'
```

Retrieve SVMs configured with audit log volumes

GET /storage/snaplock/audit-logs

Introduced In: 9.7

Retrieves a list of SVMs configured with audit log volumes.

Related ONTAP commands

- `snaplock log show`

Learn more

- [DOC /storage/snaplock/audit-logs](#)

Parameters

Name	Type	In	Required	Description
log_volume.retention_period	string	query	False	Filter by log_volume.retention_period
log_volume.volume.uuid	string	query	False	Filter by log_volume.volume.uuid
log_volume.volume.name	string	query	False	Filter by log_volume.volume.name
log_volume.max_log_size	integer	query	False	Filter by log_volume.max_log_size
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
log_files.base_name	string	query	False	Filter by log_files.base_name
log_files.expiry_time	string	query	False	Filter by log_files.expiry_time
log_files.path	string	query	False	Filter by log_files.path
log_files.size	integer	query	False	Filter by log_files.size
log_archive.base_name	string	query	False	Filter by log_archive.base_name

Name	Type	In	Required	Description
log_archive.expiry_time	string	query	False	Filter by log_archive.expiry_time
log_archive.path	string	query	False	Filter by log_archive.path
log_archive.size	integer	query	False	Filter by log_archive.size
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[snaplock_log]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "log_archive": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "base_name": "string",
        "expiry_time": "2058-06-04T19:00:00Z",
        "path": "/snaplock_log/system_logs/20180822_005947_GMT-
present",
        "size": "20000"
      },
      "log_files": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "base_name": "string",
          "expiry_time": "2058-06-04T19:00:00Z",
          "path": "/snaplock_log/system_logs/20180822_005947_GMT-
present",
          "size": "20000"
        }
      ],
      "log_volume": {
        "_links": {
```

```

        "self": {
            "href": "/api/resourcelink"
        },
        "max_log_size": "20971520",
        "retention_period": "P30M",
        "volume": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        },
        "svm": {
            "_links": {
                "self": {
                    "href": "/api/resourcelink"
                }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        }
    ]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

log_archive

Name	Type	Description
_links	_links	
archive	boolean	Archive the specified SnapLock log file for the given base_name, and create a new log file. If base_name is not mentioned, archive all log files.
base_name	string	Base name of log file
expiry_time	string	Expiry time of the log file in date-time format. Value '9999-12-31T00:00:00Z' indicates infinite expiry time.
path	string	Absolute path of the log file in the volume
size	integer	Size of the log file in bytes

snaplock_log_file

Name	Type	Description
_links	_links	

Name	Type	Description
archive	boolean	Archive the specified SnapLock log file for the given base_name, and create a new log file. If base_name is not mentioned, archive all log files.
base_name	string	Base name of log file
expiry_time	string	Expiry time of the log file in date-time format. Value '9999-12-31T00:00:00Z' indicates infinite expiry time.
path	string	Absolute path of the log file in the volume
size	integer	Size of the log file in bytes

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

snaplock_log_volume

Name	Type	Description
_links	_links	
max_log_size	integer	Maximum size of log file in bytes

Name	Type	Description
retention_period	string	Specifies the default log record retention period. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, minutes and seconds. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. A duration in hours, minutes and seconds is represented by "PT<num>H", "PT<num>M", and "PT<num>S" respectively. The period string must contain only a single time element i.e. either years, months, days, hours, minutes or seconds. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the string "infinite".</num></num></num></num></num></num>
volume	volume	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snaplock_log

Name	Type	Description
_links	_links	
log_archive	log_archive	

Name	Type	Description
log_files	array[snaplock_log_file]	
log_volume	snaplock_log_volume	
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a SnapLock log configuration for an SVM

POST /storage/snaplock/audit-logs

Introduced In: 9.7

Creates a SnapLock log configuration for an SVM. A SnapLock log configuration consists of a volume to store the log, the maximum size of the log file, and the default period of time for which the log file should be retained. The input parameter `retention_period` expects the duration in ISO 8601 format.

Required properties

- `svm.uuid` or `svm.name` - Name or UUID of the SVM.
- `log_volume.volume.name` or `log_volume.volume.uuid` - Name or UUID of audit log volume.

Recommended optional properties

- `log_volume.max_log_size` - Max log file size.
- `log_volume.volume.retention_period` - Retention period of log file.

Default property values

If not specified in POST, the following default property values are assigned:

- `log_volume.retention_period` - *P6M*
- `log_volume.max_log_size` - *10MB*

Related ONTAP commands

- `snaplock log create`

Learn more

- [DOC /storage/snaplock/audit-logs](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
log_archive	log_archive	
log_files	array[snaplock_log_file]	
log_volume	snaplock_log_volume	
svm	svm	

Example request

```
{
  "log_archive": {
    "base_name": "string",
    "expiry_time": "2058-06-04T19:00:00Z",
    "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
    "size": "20000"
  },
  "log_files": [
    {
      "base_name": "string",
      "expiry_time": "2058-06-04T19:00:00Z",
      "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
      "size": "20000"
    }
  ],
  "log_volume": {
    "max_log_size": "20971520",
    "retention_period": "P30M",
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090340	{field} is a required field
14090343	Invalid Field
14090346	Internal Error. Wait a few minutes, then try the command again

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

log_archive

Name	Type	Description
archive	boolean	Archive the specified SnapLock log file for the given base_name, and create a new log file. If base_name is not mentioned, archive all log files.
base_name	string	Base name of log file
expiry_time	string	Expiry time of the log file in date-time format. Value '9999-12-31T00:00:00Z' indicates infinite expiry time.
path	string	Absolute path of the log file in the volume
size	integer	Size of the log file in bytes

snaplock_log_file

Name	Type	Description
archive	boolean	Archive the specified SnapLock log file for the given base_name, and create a new log file. If base_name is not mentioned, archive all log files.
base_name	string	Base name of log file
expiry_time	string	Expiry time of the log file in date-time format. Value '9999-12-31T00:00:00Z' indicates infinite expiry time.

Name	Type	Description
path	string	Absolute path of the log file in the volume
size	integer	Size of the log file in bytes

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

snaplock_log_volume

Name	Type	Description
max_log_size	integer	Maximum size of log file in bytes

Name	Type	Description
retention_period	string	Specifies the default log record retention period. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, minutes and seconds. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. A duration in hours, minutes and seconds is represented by "PT<num>H", "PT<num>M", and "PT<num>S" respectively. The period string must contain only a single time element i.e. either years, months, days, hours, minutes or seconds. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the string "infinite".</num></num></num></num></num></num>
volume	volume	

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snaplock_log

Name	Type	Description
log_archive	log_archive	
log_files	array[snaplock_log_file]	
log_volume	snaplock_log_volume	

Name	Type	Description
svm	svm	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Disassociate SnapLock audit logs

```
DELETE /storage/snaplock/audit-logs/{svm.uuid}
```

Introduced In: 9.7

Disassociates a SnapLock volume as the audit log volume for an SVM. This API closes all the active log files in the log volume and marks the volume as disabled for SnapLock logging.

Related ONTAP commands

- `snaplock log delete`

Learn more

- [DOC /storage/snaplock/audit-logs](#)

Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	SVM UUID
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090346	Internal Error. Wait a few minutes, then try the command again

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve SnapLock audit logs

GET /storage/snaplock/audit-logs/{svm.uuid}

Introduced In: 9.7

Retrieves SnapLock logs for the specified SVM.

Related ONTAP commands

- `snaplock log show`

Learn more

- [DOC /storage/snaplock/audit-logs](#)

Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	SVM UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
log_archive	log_archive	
log_files	array[snaplock_log_file]	
log_volume	snaplock_log_volume	
svm	svm	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "log_archive": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "base_name": "string",
    "expiry_time": "2058-06-04T19:00:00Z",
    "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
    "size": "20000"
  },
  "log_files": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "base_name": "string",
      "expiry_time": "2058-06-04T19:00:00Z",
      "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
      "size": "20000"
    }
  ],
  "log_volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_log_size": "20971520",
    "retention_period": "P30M",
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      }
    }
  },
}
```

```

    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

log_archive

Name	Type	Description
_links	_links	
archive	boolean	Archive the specified SnapLock log file for the given base_name, and create a new log file. If base_name is not mentioned, archive all log files.
base_name	string	Base name of log file
expiry_time	string	Expiry time of the log file in date-time format. Value '9999-12-31T00:00:00Z' indicates infinite expiry time.
path	string	Absolute path of the log file in the volume
size	integer	Size of the log file in bytes

snaplock_log_file

Name	Type	Description
_links	_links	
archive	boolean	Archive the specified SnapLock log file for the given base_name, and create a new log file. If base_name is not mentioned, archive all log files.
base_name	string	Base name of log file

Name	Type	Description
expiry_time	string	Expiry time of the log file in date-time format. Value '9999-12-31T00:00:00Z' indicates infinite expiry time.
path	string	Absolute path of the log file in the volume
size	integer	Size of the log file in bytes

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

snaplock_log_volume

Name	Type	Description
_links	_links	
max_log_size	integer	Maximum size of log file in bytes

Name	Type	Description
retention_period	string	Specifies the default log record retention period. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, minutes and seconds. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. A duration in hours, minutes and seconds is represented by "PT<num>H", "PT<num>M", and "PT<num>S" respectively. The period string must contain only a single time element i.e. either years, months, days, hours, minutes or seconds. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the string "infinite".</num></num></num></num></num></num>
volume	volume	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update SnapLock audit logs

PATCH /storage/snaplock/audit-logs/{svm.uuid}

Introduced In: 9.7

Updates one of the following:

- the audit log volume,
- the attributes of the audit log volume present, or
- archive the current audit log files

Related ONTAP commands

- `snaplock log modify`

Learn more

- [DOC /storage/snaplock/audit-logs](#)

Parameters

Name	Type	In	Required	Description
svm.uuid	string	path	True	SVM UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
log_archive	log_archive	
log_volume	snaplock_log_volume	

Example request

```
{
  "log_archive": {
    "base_name": "string",
    "expiry_time": "2058-06-04T19:00:00Z",
    "path": "/snaplock_log/system_logs/20180822_005947_GMT-present",
    "size": "20000"
  },
  "log_volume": {
    "max_log_size": "20971520",
    "retention_period": "P30M",
    "volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090344	If log_volume is specified, then log_archive must not be specified
14090345	If log_archive.base_name is specified, then log_archive.archive must also be specified
14090346	Internal Error. Wait a few minutes, then try the command again

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

log_archive

Name	Type	Description
archive	boolean	Archive the specified SnapLock log file for the given base_name, and create a new log file. If base_name is not mentioned, archive all log files.
base_name	string	Base name of log file
expiry_time	string	Expiry time of the log file in date-time format. Value '9999-12-31T00:00:00Z' indicates infinite expiry time.
path	string	Absolute path of the log file in the volume
size	integer	Size of the log file in bytes

snaplock_log_file

Name	Type	Description
archive	boolean	Archive the specified SnapLock log file for the given base_name, and create a new log file. If base_name is not mentioned, archive all log files.
base_name	string	Base name of log file
expiry_time	string	Expiry time of the log file in date-time format. Value '9999-12-31T00:00:00Z' indicates infinite expiry time.

Name	Type	Description
path	string	Absolute path of the log file in the volume
size	integer	Size of the log file in bytes

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

snaplock_log_volume

Name	Type	Description
max_log_size	integer	Maximum size of log file in bytes

Name	Type	Description
retention_period	string	Specifies the default log record retention period. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, minutes and seconds. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. A duration in hours, minutes and seconds is represented by "PT<num>H", "PT<num>M", and "PT<num>S" respectively. The period string must contain only a single time element i.e. either years, months, days, hours, minutes or seconds. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the string "infinite".</num></num></num></num></num></num>
volume	volume	

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snaplock_log

Name	Type	Description
log_archive	log_archive	
log_volume	snaplock_log_volume	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage the SnapLock ComplianceClock

Storage SnapLock compliance-clocks endpoint overview

This API manages the ComplianceClock of the system. ComplianceClock determines the expiry time of the SnapLock objects in the system. The user can initialize the ComplianceClock once and when it is set, it cannot be changed by the user. ComplianceClock initialize is not supported in REST.

Retrieve the SnapLock ComplianceClock for nodes in a cluster

GET /storage/snaplock/compliance-clocks

Introduced In: 9.7

Retrieves the SnapLock ComplianceClock for all of the nodes in the cluster.

Related ONTAP commands

- `snaplock compliance-clock show`

Learn more

- [DOC /storage/snaplock/compliance-clocks](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none">• Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[snaplock_compliance_clock]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "time": "2018-06-04T19:00:00Z"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

snaplock_compliance_clock

Name	Type	Description
_links	_links	
node	node	
time	string	Compliance clock time

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a SnapLock ComplianceClock for a node

GET /storage/snaplock/compliance-clocks/{node.uuid}

Introduced In: 9.7

Retrieves the SnapLock ComplianceClock for a specific node.

Related ONTAP commands

- `snaplock compliance-clock show`

Learn more

- [DOC /storage/snaplock/compliance-clocks](#)

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
node	node	
time	string	Compliance clock time

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "time": "2018-06-04T19:00:00Z"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage event based retention operations

Storage SnapLock event-retention operations endpoint overview

Use this API to display all Event Based Retention (EBR) operations and to apply an EBR policy on a specified

volume.

Examples

1. Displays all of the EBR operations:

```
GET "/api/storage/snaplock/event-retention/operations"
```

1. Displays all completed EBR operations:

```
GET "/api/storage/snaplock/event-retention/operations?state=completed"
```

1. Displays all completed EBR operations with filter set as volume.uuid:

```
GET "/api/storage/snaplock/event-retention/operations?volume.uuid=b96f976e-404b-11e9-bff2-0050568e4dbe"
```

1. Displays all of the EBR operations with filter set as volume.name:

```
GET "/api/storage/snaplock/event-retention/operations?volume.name=SLCVOL"
```

Examples

1. Applies an EBR policy on a specific path:

```
POST "/api/storage/snaplock/event-retention/operations"
'{"volume.name":"SLCVOL", "policy.name":"p1day", "path":"/dir1/file.txt"}'
```

1. Applies an EBR policy on the complete volume:

```
POST "/api/storage/snaplock/event-retention/operations"
'{"volume.name":"SLCVOL", "policy.name":"p1day", "path":"/"}'
```

Example

```
DELETE "/api/storage/snaplock/event-retention/operations/16842999"
```

Retrieve EBR operations

```
GET /storage/snaplock/event-retention/operations
```

Introduced In: 9.7

Retrieves a list of all EBR operations.

Related ONTAP commands

- `snaplock event-retention show`

Learn more

- [DOC /storage/snaplock/event-retention/operations](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	query	True	Volume UUID
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[ebr_operation]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": "16842759",
      "num_files_failed": "0",
      "num_files_processed": "50",
      "num_files_skipped": "2",
      "num_inodes_ignored": "2",
      "path": "/dir1/file",
      "policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "name": "string",
      "retention_period": "P30M",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    {
      "state": "completed",
      "svm": {
        "_links": {
          "self": {
```

```

        "href": "/api/resourcelink"
      },
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```


See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

policy

Name	Type	Description
_links	_links	
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".
svm	svm	SVM, applies only to SVM-scoped objects.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

ebr_operation

Name	Type	Description
_links	_links	
id	integer	Operation ID
num_files_failed	integer	Specifies the number of files on which the application of EBR policy failed.
num_files_processed	integer	Specifies the number of files on which EBR policy was applied successfully.
num_files_skipped	integer	Specifies the number of files on which the application of EBR policy was skipped.
num_inodes_ignored	integer	Specifies the number of inodes on which the application of EBR policy was not attempted because they were not regular files.
path	string	The path for the EBR operation. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
policy	policy	
state	string	Specifies the operation status of an EBR operation.

Name	Type	Description
svm	svm	
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create an EBR policy

POST /storage/snaplock/event-retention/operations

Introduced In: 9.7

Creates an EBR policy.

Required properties

- `path` - Path of the file.
- `policy.name` - Name of the EBR policy.

Related ONTAP commands

- `snaplock event-retention apply`

Learn more

- [DOC /storage/snaplock/event-retention/operations](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
id	integer	Operation ID
num_files_failed	integer	Specifies the number of files on which the application of EBR policy failed.
num_files_processed	integer	Specifies the number of files on which EBR policy was applied successfully.
num_files_skipped	integer	Specifies the number of files on which the application of EBR policy was skipped.
num_inodes_ignored	integer	Specifies the number of inodes on which the application of EBR policy was not attempted because they were not regular files.
path	string	The path for the EBR operation. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
policy	policy	
state	string	Specifies the operation status of an EBR operation.
svm	svm	
volume	volume	

Example request

```
{
  "id": "16842759",
  "num_files_failed": "0",
  "num_files_processed": "50",
  "num_files_skipped": "2",
  "num_inodes_ignored": "2",
  "path": "/dir1/file",
  "policy": {
    "name": "string",
    "retention_period": "P30M",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "completed",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 201, Created

Name	Type	Description
id	integer	Operation ID
num_files_failed	integer	Specifies the number of files on which the application of EBR policy failed.
num_files_processed	integer	Specifies the number of files on which EBR policy was applied successfully.

Name	Type	Description
num_files_skipped	integer	Specifies the number of files on which the application of EBR policy was skipped.
num_inodes_ignored	integer	Specifies the number of inodes on which the application of EBR policy was not attempted because they were not regular files.
path	string	The path for the EBR operation. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
policy	policy	
state	string	Specifies the operation status of an EBR operation.
svm	svm	
volume	volume	

Example response

```
{
  "id": "16842759",
  "num_files_failed": "0",
  "num_files_processed": "50",
  "num_files_skipped": "2",
  "num_inodes_ignored": "2",
  "path": "/dir1/file",
  "policy": {
    "name": "string",
    "retention_period": "P30M",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "state": "completed",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

policy

Name	Type	Description
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".
svm	svm	SVM, applies only to SVM-scoped objects.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

ebr_operation

Name	Type	Description
id	integer	Operation ID
num_files_failed	integer	Specifies the number of files on which the application of EBR policy failed.
num_files_processed	integer	Specifies the number of files on which EBR policy was applied successfully.
num_files_skipped	integer	Specifies the number of files on which the application of EBR policy was skipped.
num_inodes_ignored	integer	Specifies the number of inodes on which the application of EBR policy was not attempted because they were not regular files.
path	string	The path for the EBR operation. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
policy	policy	
state	string	Specifies the operation status of an EBR operation.
svm	svm	
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Cancel an ongoing EBR operation

DELETE /storage/snaplock/event-retention/operations/{id}

Introduced In: 9.7

Aborts an ongoing EBR operation.

Related ONTAP commands

- `snaplock event-retention abort`

Learn more

- [DOC /storage/snaplock/event-retention/operations](#)

Parameters

Name	Type	In	Required	Description
id	string	path	True	Unique identifier of an EBR operation

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090541	A completed or failed operation cannot be aborted

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve attributes for an EBR operation

GET /storage/snaplock/event-retention/operations/{id}

Introduced In: 9.7

Retrieves a list of attributes for an EBR operation.

Related ONTAP commands

- `snaplock event-retention show`

Learn more

- [DOC /storage/snaplock/event-retention/operations](#)

Parameters

Name	Type	In	Required	Description
id	string	path	True	Unique identifier of an EBR operation
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
id	integer	Operation ID
num_files_failed	integer	Specifies the number of files on which the application of EBR policy failed.
num_files_processed	integer	Specifies the number of files on which EBR policy was applied successfully.
num_files_skipped	integer	Specifies the number of files on which the application of EBR policy was skipped.
num_inodes_ignored	integer	Specifies the number of inodes on which the application of EBR policy was not attempted because they were not regular files.
path	string	The path for the EBR operation. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
policy	policy	
state	string	Specifies the operation status of an EBR operation.
svm	svm	
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": "16842759",
  "num_files_failed": "0",
  "num_files_processed": "50",
  "num_files_skipped": "2",
  "num_inodes_ignored": "2",
  "path": "/dir1/file",
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "string",
  "retention_period": "P30M",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"state": "completed",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
},
"name": "svm1",
"uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}
```

```
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

policy

Name	Type	Description
_links	_links	
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".
svm	svm	SVM, applies only to SVM-scoped objects.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage event based retention policies

Storage SnapLock event-retention policies endpoint overview

An event retention policy consists of a policy-name and a retention-period. The policy can be applied to a single file or files in a directory. Only a user with the security login role vsadmin-snaplock can perform the operation. EBR policies cannot be applied to files under a Legal-Hold.

Examples

1. Creates an EBR policy policy_name with a retention period of "10 years":

```
POST "/api/storage/snaplock/event-retention/policies/" '{"name":  
"policy_name","retention_period": "P10Y"}'
```

1. Creates an EBR policy policy_name1 with a retention period of "infinite":

```
POST "/api/storage/snaplock/event-retention/policies/" '{"name":  
"policy_name1","retention_period": "infinite"}'
```

Retrieve event retention policies for an SVM

GET /storage/snaplock/event-retention/policies

Introduced In: 9.7

Retrieves all event retention policies for an SVM.

Related ONTAP commands

- `snaplock event-retention policy show`

Learn more

- [DOC /storage/snaplock/event-retention/policies](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[snaplock_retention_policy]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "string",
      "retention_period": "P30M",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snaplock_retention_policy

Name	Type	Description
_links	_links	
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".</num></num></num></num></num>
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Create an EBR policy for an SVM

POST /storage/snaplock/event-retention/policies

Introduced In: 9.7

Creates an Event Based Retention (EBR) policy for an SVM. The input parameter `retention_period` expects the duration in ISO 8601 format or infinite.

Required properties

- `name` - Event retention policy name.
- `retention_period` - Retention period of the EBR policy.

Related ONTAP commands

- `snaplock event-retention policy create`

Learn more

- [DOC /storage/snaplock/event-retention/policies](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".</num></num></num></num></num>
svm	svm	

Example request

```
{
  "name": "string",
  "retention_period": "P30M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 201, Created

Name	Type	Description
name	string	Specifies the EBR policy name
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".</num></num></num></num></num>
svm	svm	

Example response

```
{
  "name": "string",
  "retention_period": "P30M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snaplock_retention_policy

Name	Type	Description
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".</num></num></num></num></num>
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Delete an EBR policy

DELETE /storage/snaplock/event-retention/policies/{policy.name}

Introduced In: 9.7

Deletes the specified Event Based Retention (EBR) policy.

Related ONTAP commands

- `snaplock event-retention policy delete`

Learn more

- [DOC /storage/snaplock/event-retention/policies](#)

Parameters

Name	Type	In	Required	Description
policy.name	string	path	True	Name of the retention policy

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve attributes of an EBR policy

GET /storage/snaplock/event-retention/policies/{policy.name}

Introduced In: 9.7

Retrieves a list of attributes of the specified Event Based Retention (EBR) policy.

Related ONTAP commands

- `snaplock event-retention policy show`

Learn more

- [DOC /storage/snaplock/event-retention/policies](#)

Parameters

Name	Type	In	Required	Description
policy.name	string	path	True	Name of the retention policy
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".</num></num></num></num></num>
svm	svm	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "string",
  "retention_period": "P30M",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the retention period of an EBR policy

PATCH /storage/snaplock/event-retention/policies/{policy.name}

Introduced In: 9.7

Updates the retention period of an Event Based Retention (EBR) policy.

Related ONTAP commands

- `snaplock event-retention policy modify`

Example

Updates the retention period of an EBR policy "policy_name":

```
PATCH "/api/storage/snaplock/event-retention/policies/" '{"name":  
"policy_name","retention_period": "P20Y"}'
```

Learn more

- [DOC /storage/snaplock/event-retention/policies](#)

Parameters

Name	Type	In	Required	Description
policy.name	string	path	True	Name of the retention policy

Request Body

Name	Type	Description
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".</num></num></num></num></num>
svm	svm	

Example request

```
{
  "name": "string",
  "retention_period": "P30M",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snaplock_retention_policy

Name	Type	Description
name	string	Specifies the EBR policy name

Name	Type	Description
retention_period	string	Specifies the retention period of an event based retention policy. The retention period value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours or minutes. A period specified for years, months and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively. For example "P10Y" represents a duration of 10 years. Similarly, a duration in hours, minutes is represented by "PT<num>H", "PT<num>M" respectively. The period string must contain only a single time element i.e. either years, months, days, hours or minutes. A duration which combines different periods is not supported, example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the retention period field also accepts the strings "infinite" and "unspecified".</num></num></num></num></num>
svm	svm	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

View key information about files and volumes

Storage SnapLock file-fingerprints endpoint overview

Use this API to view key information about files and volumes, including the file type (regular, WORM, or WORM appendable), the volume expiration date, and so on.

Retrieve fingerprint operations of an SVM and volume

GET /storage/snaplock/file-fingerprints

Introduced In: 9.7

Retrieves a list of all the fingerprint operations of the specified SVM and volume.

Related ONTAP commands

- `volume file fingerprint show`

Example

```
GET "/api/storage/snaplock/file-fingerprints/?svm.uuid=23940494-3f3a-11e9-8675-0050568e8f89&volume.uuid=36cdb58c-3f3a-11e9-8675-0050568e8f89"
```

Learn more

- [DOC /storage/snaplock/file-fingerprints](#)

Parameters

Name	Type	In	Required	Description
svm.uuid	string	query	True	SVM UUID
volume.uuid	string	query	True	Volume UUID
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of Records.
records	array[snaplock_file_fingerprint]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "algorithm": "sha256",
      "data_fingerprint":
"MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=",
      "file_size": "1048576",
      "file_type": "worm",
      "id": "17039367",
      "metadata_fingerprint":
"8iMjqJXiNcggXT5XuRhLiEwIrJEihDmws0hrexnjgmc=",
      "path": "/homedir/dir1",
      "scope": "data_only",
      "state": "completed",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

```
}
}
]
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090442	Invalid Vserver UUID
14090445	Invalid key values. Provide Vserver UUID and Volume UUID

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

snaplock_file_fingerprint

Name	Type	Description
_links	_links	
algorithm	string	The digest algorithm which is used for the fingerprint computation
data_fingerprint	string	The digest value of data of the file. The fingerprint is base64 encoded. This field is not included if the scope is metadata-only.
file_size	integer	The size of the file in bytes.
file_type	string	The type of the file.
id	integer	A unique identifier for the fingerprint operation
metadata_fingerprint	string	The digest value of metadata of the file. The metadata fingerprint is calculated for file size, file ctime, file mtime, file crtime, file retention time, file uid, file gid, and file type. The fingerprint is base64 encoded. This field is not included if the scope is data-only.
path	string	Specifies the path on which file fingerprint operation is running or has completed. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
scope	string	The scope of the file which is used for the fingerprint computation
state	string	Specifies the status of fingerprint operation.
svm	svm	
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a fingerprint computation session

POST /storage/snaplock/file-fingerprints

Introduced In: 9.7

Creates a fingerprint computation session on the file and returns a session-id. This session-id is a unique identifier that you can use to retrieve the progress of an ongoing fingerprint operation. When the operation is complete, you can use the session-id to retrieve the complete fingerprint output for the file .

Required properties

- `svm.uuid` or `svm.name` - Name or UUID of the SVM.
- `volume.name` or `volume.uuid` - Name or UUID of the volume.
- `path` - Path of the file.

Default property values

If not specified in POST, the follow default property values are assigned:

- `algorithm` - *md5*

Related ONTAP commands

- `volume file fingerprint start`

Example

```
POST "/api/storage/snaplock/file-fingerprints" '{"svm":{"uuid":"23940494-3f3a-11e9-8675-0050568e8f89"},"volume": {"uuid":"26cdb58c-3f3a-11e9-8675-0050568e8f89"},"path":"/vol/a1.txt","algorithm":"md5"}'
```

Learn more

- [DOC /storage/snaplock/file-fingerprints](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
algorithm	string	The digest algorithm which is used for the fingerprint computation
data_fingerprint	string	The digest value of data of the file. The fingerprint is base64 encoded. This field is not included if the scope is metadata-only.
file_size	integer	The size of the file in bytes.
file_type	string	The type of the file.
id	integer	A unique identifier for the fingerprint operation
metadata_fingerprint	string	The digest value of metadata of the file. The metadata fingerprint is calculated for file size, file ctime, file mtime, file crtime, file retention time, file uid, file gid, and file type. The fingerprint is base64 encoded. This field is not included if the scope is data-only.
path	string	Specifies the path on which file fingerprint operation is running or has completed. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
scope	string	The scope of the file which is used for the fingerprint computation
state	string	Specifies the status of fingerprint operation.

Name	Type	Description
svm	svm	
volume	volume	

Example request

```
{
  "algorithm": "sha256",
  "data_fingerprint": "MOFJVeVxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=",
  "file_size": "1048576",
  "file_type": "worm",
  "id": "17039367",
  "metadata_fingerprint":
"8iMjqJXiNcggXT5XuRhLiEwIrJEihDmwS0hrexnjgmc=",
  "path": "/homedir/dir1",
  "scope": "data_only",
  "state": "completed",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090443	Invalid volume name
14090444	Invalid Vserver name
14090447	Invalid volume UUID
14090448	Invalid key values. Provide valid Vserver name and volume name or Vserver UUID and volume UUID

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

snaplock_file_fingerprint

Name	Type	Description
algorithm	string	The digest algorithm which is used for the fingerprint computation
data_fingerprint	string	The digest value of data of the file. The fingerprint is base64 encoded. This field is not included if the scope is metadata-only.
file_size	integer	The size of the file in bytes.

Name	Type	Description
file_type	string	The type of the file.
id	integer	A unique identifier for the fingerprint operation
metadata_fingerprint	string	The digest value of metadata of the file. The metadata fingerprint is calculated for file size, file ctime, file mtime, file crtime, file retention time, file uid, file gid, and file type. The fingerprint is base64 encoded. This field is not included if the scope is data-only.
path	string	Specifies the path on which file fingerprint operation is running or has completed. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
scope	string	The scope of the file which is used for the fingerprint computation
state	string	Specifies the status of fingerprint operation.
svm	svm	
volume	volume	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Cancel a fingerprint operation

```
DELETE /storage/snaplock/file-fingerprints/{id}
```

Introduced In: 9.7

Aborts an in-progress fingerprint operation. This API takes session-id as input and aborts the fingerprint operation that is associated with the specified session-id.

Related ONTAP commands

- `volume file fingerprint abort`

Learn more

- [DOC /storage/snaplock/file-fingerprints](#)

Parameters

Name	Type	In	Required	Description
id	string	path	True	Unique identifier of the operation

Response

```
Status: 200, Ok
```

Error

```
Status: Default
```

ONTAP Error Response codes

Error code	Description
14090440	File fingerprint operation has completed
14090446	Invalid session ID

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the file fingerprint information for a specific session ID

GET /storage/snaplock/file-fingerprints/{id}

Introduced In: 9.7

Retrieves the file fingerprint information for a specific session ID.

Related ONTAP commands

- `volume file fingerprint dump`

Learn more

- [DOC /storage/snaplock/file-fingerprints](#)

Parameters

Name	Type	In	Required	Description
id	string	path	True	Unique identifier of the operation
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
algorithm	string	The digest algorithm which is used for the fingerprint computation
data_fingerprint	string	The digest value of data of the file. The fingerprint is base64 encoded. This field is not included if the scope is metadata-only.
file_size	integer	The size of the file in bytes.
file_type	string	The type of the file.
id	integer	A unique identifier for the fingerprint operation
metadata_fingerprint	string	The digest value of metadata of the file. The metadata fingerprint is calculated for file size, file ctime, file mtime, file crtime, file retention time, file uid, file gid, and file type. The fingerprint is base64 encoded. This field is not included if the scope is data-only.
path	string	Specifies the path on which file fingerprint operation is running or has completed. Specifies the path relative to the output volume root, of the form "/path". The path can be path to a file or a directory.
scope	string	The scope of the file which is used for the fingerprint computation
state	string	Specifies the status of fingerprint operation.
svm	svm	
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "algorithm": "sha256",
  "data_fingerprint": "MOFJVevxNSJm3C/4Bn5oEEYH51CrudOzZYK4r5Cfy1g=",
  "file_size": "1048576",
  "file_type": "worm",
  "id": "17039367",
  "metadata_fingerprint":
"8iMjqJXiNcqqXT5XuRhLiEwIrJEihDmwS0hrexnjgmc=",
  "path": "/homedir/dir1",
  "scope": "data_only",
  "state": "completed",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090441	Invalid key values. Provide session ID or Vserver UUID and volume UUID
14090449	File fingerprint operation has not completed

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage SnapLock file retention time

Storage SnapLock file volume.uuid path endpoint overview

This API manages the SnapLock retention time of a file. You can perform a privileged-delete operation by executing this API.

Examples

1. Sets the SnapLock retention time of a file:

```
PATCH "/api/storage/snaplock/file/000dc5fd-4175-11e9-b937-0050568e3f82/%2Ffile2.txt" '{"expiry_time": "2030-02-14T18:30:00+5:30"}'
```

2. Extends the retention time of a WORM file:

```
PATCH "/api/storage/snaplock/file/000dc5fd-4175-11e9-b937-0050568e3f82/%2Ffile2.txt" '{"expiry_time": "infinite"}'
```

Delete unexpired WORM files of a SnapLock enterprise volume

DELETE /storage/snaplock/file/{volume.uuid}/{path}

Introduced In: 9.7

Deletes unexpired WORM files of a SnapLock Enterprise volume. This is a privileged-delete operation. The only built-in role that has access to the command is vsadmin-snaplock.

Related ONTAP commands

- `volume file privileged-delete`

Learn more

- [DOC /storage/snaplock/file/{volume.uuid}/{path}](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
path	string	path	True	Path of the file in the form "/<dirpath>/<filename>" </filename></dirpath>
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090347	File path must be in the format "<dir><file path></file></dir>"

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve SnapLock retention details of a file

GET /storage/snaplock/file/{volume.uuid}/{path}

Introduced In: 9.7

Retrieves the SnapLock retention details of the specified file. An indefinite expiry time indicates the file is under a Legal-Hold.

Related ONTAP commands

- `volume file retention show`

Learn more

- [DOC /storage/snaplock/file/{volume.uuid}/{path}](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
path	string	path	True	Path of the file in the form "/<dirpath>/<filename>" </filename></dirpath>
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
expiry_time	string	Expiry time of the file in date-time format, "infinite", "indefinite", or "unspecified". An "infinite" retention time indicates that the file will be retained forever. An "unspecified" retention time indicates that the file will be retained forever; however, the retention time of the file can be changed to an absolute value. An "indefinite" retention time indicates that the file is under Legal-Hold.

Name	Type	Description
file_path	string	Specifies the volume relative path of the file
svm	svm	
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "expiry_time": "2058-06-04T19:00:00Z",
  "file_path": "/dir1/file",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090347	File path must be in the format "<dir><file path></file></dir>"

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the SnapLock retention time of a file

PATCH /storage/snaplock/file/{volume.uuid}/{path}

Introduced In: 9.7

Updates the SnapLock retention time of a file or extends the retention time of a WORM file. Input parameter "expiry_time" expects the date in ISO 8601 format, "infinite", or "unspecified".

Related ONTAP commands

- `volume file retention set`

Learn more

- [DOC /storage/snaplock/file/{volume.uuid}/{path}](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
path	string	path	True	Path of the file in the form "/<dirpath>/<filename>" </filename></dirpath>

Request Body

Name	Type	Description
expiry_time	string	Expiry time of the file in date-time format, "infinite", "indefinite", or "unspecified". An "infinite" retention time indicates that the file will be retained forever. An "unspecified" retention time indicates that the file will be retained forever; however, the retention time of the file can be changed to an absolute value. An "indefinite" retention time indicates that the file is under Legal-Hold.

Example request

```
{
  "expiry_time": "2058-06-04T19:00:00Z"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090347	File path must be in the format "<dir><file path></file></dir>"
14090348	Invalid Expiry time

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

snaplock_file_retention

Name	Type	Description
expiry_time	string	Expiry time of the file in date-time format, "infinite", "indefinite", or "unspecified". An "infinite" retention time indicates that the file will be retained forever. An "unspecified" retention time indicates that the file will be retained forever; however, the retention time of the file can be changed to an absolute value. An "indefinite" retention time indicates that the file is under Legal-Hold.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage SnapLock storage litigations

Storage SnapLock litigations endpoint overview

Use this API to retain Compliance-mode WORM files for the duration of a litigation. A file under a legal-hold behaves as a WORM file with an indefinite retention period. Litigation ID is a combination of volume UUID and litigation name in the format <volume UUID>:<litigation name>. Only a user with the security login role vsadmin-snaplock can perform the operation.

Retrieve litigations under an SVM

GET /storage/snaplock/litigations

Introduced In: 9.7

Retrieves the list of litigations under an SVM.

Related ONTAP commands

- `snaplock legal-hold show`

Learn more

- [DOC /storage/snaplock/litigations](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of Records
records	array[snaplock_litigation]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": "string",
      "name": "lit1",
      "operations": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "id": "16842759",
          "num_files_failed": "0",
          "num_files_processed": "30",
          "num_files_skipped": "10",
          "num_inodes_ignored": "10",
          "path": "/dir1",
          "state": "completed",
          "type": "begin"
        }
      ],
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ],
}
```

```

    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ]
}

```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090346	Internal Error. Wait a few minutes, then try the command again

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```


See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

snaplock_legal_hold_operation

Name	Type	Description
_links	_links	
id	integer	Operation ID.
num_files_failed	string	Specifies the number of files on which legal-hold operation failed.
num_files_processed	string	Specifies the number of files on which legal-hold operation was successful.
num_files_skipped	string	Specifies the number of files on which legal-hold begin operation was skipped. The legal-hold begin operation is skipped on a file if it is already under hold for a given litigation.
num_inodes_ignored	string	Specifies the number of inodes on which the legal-hold operation was not attempted because they were not regular files.
path	string	Specifies the path on which legal-hold operation is applied.

Name	Type	Description
state	string	Specifies the status of legal-hold operation.
type	string	Specifies the type of legal-hold operation.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

snaplock_litigation

Name	Type	Description
_links	_links	
id	string	Specifies the litigation ID.
name	string	Specifies the legal-hold litigation name.
operations	array[snaplock_legal_hold_operation]	
svm	svm	

Name	Type	Description
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Start a legal-hold operation

POST /storage/snaplock/litigations

Introduced In: 9.7

Starts a Legal-Hold.

Required properties

- `path` - Path of the file.
- `name` - Litigation name.
- `volume.name` or `volume.uuid` - Name or UUID of the volume.

Related ONTAP commands

- `snaplock legal-hold begin`

Example

```
POST "/api/storage/snaplock/litigations"
'{ "volume.name": "SLC1", "name": "l3", "path": "/b.txt" }'
```

Learn more

- [DOC /storage/snaplock/litigations](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
id	string	Specifies the litigation ID.
name	string	Specifies the legal-hold litigation name.
operations	array[snaplock_legal_hold_operation]	
path	string	Specifies the path on which legal-hold operation has to be applied.
svm	svm	
volume	volume	

Example request

```
{
  "id": "string",
  "name": "lit1",
  "operations": [
    {
      "id": "16842759",
      "num_files_failed": "0",
      "num_files_processed": "30",
      "num_files_skipped": "10",
      "num_inodes_ignored": "10",
      "path": "/dir1",
      "state": "completed"
    }
  ],
  "path": "/dir1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 201, Created

Name	Type	Description
id	string	Specifies the litigation ID.
name	string	Specifies the legal-hold litigation name.
operations	array[snaplock_legal_hold_operation]	
path	string	Specifies the path on which legal-hold operation has to be applied.
svm	svm	

Name	Type	Description
volume	volume	

Example response

```
{
  "id": "string",
  "name": "lit1",
  "operations": [
    {
      "id": "16842759",
      "num_files_failed": "0",
      "num_files_processed": "30",
      "num_files_skipped": "10",
      "num_inodes_ignored": "10",
      "path": "/dir1",
      "state": "completed"
    }
  ],
  "path": "/dir1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090346	Internal Error. Wait a few minutes, then try the command again
14090340	{field} is a required field
14090343	Invalid Field

Error code	Description
14090641	The specified volume.name and volume.uuid refer to different volumes

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

snaplock_legal_hold_operation

Name	Type	Description
id	integer	Operation ID.
num_files_failed	string	Specifies the number of files on which legal-hold operation failed.
num_files_processed	string	Specifies the number of files on which legal-hold operation was successful.
num_files_skipped	string	Specifies the number of files on which legal-hold begin operation was skipped. The legal-hold begin operation is skipped on a file if it is already under hold for a given litigation.
num_inodes_ignored	string	Specifies the number of inodes on which the legal-hold operation was not attempted because they were not regular files.
path	string	Specifies the path on which legal-hold operation is applied.
state	string	Specifies the status of legal-hold operation.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

snaplock_litigation

Name	Type	Description
id	string	Specifies the litigation ID.
name	string	Specifies the legal-hold litigation name.
operations	array[snaplock_legal_hold_operation]	
path	string	Specifies the path on which legal-hold operation has to be applied.
svm	svm	
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Create a legal-hold end on all of the files for a litigation ID

DELETE /storage/snaplock/litigations/{id}

Introduced In: 9.7

Creates a legal-hold end on all of the files for the specified litigation ID. This is only allowed when an operation is no longer in progress.

Related ONTAP commands

- `snaplock legal-hold end`

Example

```
DELETE "/api/storage/snaplock/litigations/fd72e138-4bc3-11e9-a85f-0050568eb48f%3A13"
```

Learn more

- [DOC /storage/snaplock/litigations](#)

Parameters

Name	Type	In	Required	Description
id	string	path	True	Litigation ID

Response

```
Status: 200, Ok
```

Error

```
Status: Default
```

ONTAP Error Response codes

Error code	Description
14090346	Internal Error. Wait a few minutes, then try the command again
14090343	Invalid Field
14090642	Legal-Hold operation is in progress. Legal-Hold end cannot be performed. Wait for the operation to complete or abort the operation

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve ongoing operations for a litigation ID

GET /storage/snaplock/litigations/{id}

Introduced In: 9.7

Retrieves the list of ongoing operations for the specified litigation ID.

Related ONTAP commands

- `snaplock legal-hold show`

Learn more

- [DOC /storage/snaplock/litigations](#)

Parameters

Name	Type	In	Required	Description
id	string	path	True	Litigation ID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
id	string	Specifies the litigation ID.
name	string	Specifies the legal-hold litigation name.
operations	array[snaplock_legal_hold_operation]	
svm	svm	
volume	volume	

Example response

```

{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": "string",
  "name": "lit1",
  "operations": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "id": "16842759",
      "num_files_failed": "0",
      "num_files_processed": "30",
      "num_files_skipped": "10",
      "num_inodes_ignored": "10",
      "path": "/dir1",
      "state": "completed",
      "type": "begin"
    }
  ],
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}

```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090346	Internal Error. Wait a few minutes, then try the command again
14090343	Invalid Field

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

snaplock_legal_hold_operation

Name	Type	Description
_links	_links	
id	integer	Operation ID.
num_files_failed	string	Specifies the number of files on which legal-hold operation failed.
num_files_processed	string	Specifies the number of files on which legal-hold operation was successful.
num_files_skipped	string	Specifies the number of files on which legal-hold begin operation was skipped. The legal-hold begin operation is skipped on a file if it is already under hold for a given litigation.
num_inodes_ignored	string	Specifies the number of inodes on which the legal-hold operation was not attempted because they were not regular files.
path	string	Specifies the path on which legal-hold operation is applied.
state	string	Specifies the status of legal-hold operation.
type	string	Specifies the type of legal-hold operation.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View litigation ID files

Storage SnapLock litigations litigation.id files endpoint overview

Displays the list of files under the specified litigation ID.

Display files for a litigation ID

GET /storage/snaplock/litigations/{litigation.id}/files

Introduced In: 9.7

Displays the list of files for the specified litigation ID.

Parameters

Name	Type	In	Required	Description
litigation.id	string	path	True	Litigation ID
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	List of Files under the specified litigation.

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "file": [
        "string"
      ]
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

records

Name	Type	Description
file	array[string]	Name of the file including the path from the root.
sequence_index	integer	Sequence index of files path list.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage legal-hold operations

Storage SnapLock litigations litigation.id operations endpoint overview

Manages the legal-hold operations for the specified litigation ID.

Examples

1. Adds a Legal-Hold.

```
POST "/api/storage/snaplock/litigations/f8a67b60-4461-11e9-b327-0050568ebef5:11/operations" '{"type" : "begin", "path" : "/a.txt"}'
```

1. Removes a Legal-Hold.

```
POST "/api/storage/snaplock/litigations/f8a67b60-4461-11e9-b327-0050568ebef5:11/operations" '{"type" : "end", "path" : "/a.txt"}'
```

Create or remove litigations for a path

POST /storage/snaplock/litigations/{litigation.id}/operations

Introduced In: 9.7

Creates or removes litigations for the specified path.

Required properties

- `type` - Legal-Hold operation type.
- `path` - Litigation path.

Related ONTAP commands

- `snaplock legal-hold begin`
- `snaplock legal-hold end`

Learn more

- [DOC /storage/snaplock/litigations/{litigation.id}/operations](#)

Parameters

Name	Type	In	Required	Description
litigation.id	string	path	True	Litigation ID

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
id	integer	Operation ID.
num_files_failed	string	Specifies the number of files on which legal-hold operation failed.
num_files_processed	string	Specifies the number of files on which legal-hold operation was successful.
num_files_skipped	string	Specifies the number of files on which legal-hold begin operation was skipped. The legal-hold begin operation is skipped on a file if it is already under hold for a given litigation.
num_inodes_ignored	string	Specifies the number of inodes on which the legal-hold operation was not attempted because they were not regular files.
path	string	Specifies the path on which legal-hold operation is applied.
state	string	Specifies the status of legal-hold operation.

Example request

```
{
  "id": "16842759",
  "num_files_failed": "0",
  "num_files_processed": "30",
  "num_files_skipped": "10",
  "num_inodes_ignored": "10",
  "path": "/dir1",
  "state": "completed"
}
```

Response

Status: 201, Created

Name	Type	Description
id	integer	Operation ID.
num_files_failed	string	Specifies the number of files on which legal-hold operation failed.
num_files_processed	string	Specifies the number of files on which legal-hold operation was successful.
num_files_skipped	string	Specifies the number of files on which legal-hold begin operation was skipped. The legal-hold begin operation is skipped on a file if it is already under hold for a given litigation.
num_inodes_ignored	string	Specifies the number of inodes on which the legal-hold operation was not attempted because they were not regular files.
path	string	Specifies the path on which legal-hold operation is applied.
state	string	Specifies the status of legal-hold operation.

Example response

```
{
  "id": "16842759",
  "num_files_failed": "0",
  "num_files_processed": "30",
  "num_files_skipped": "10",
  "num_inodes_ignored": "10",
  "path": "/dir1",
  "state": "completed"
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090346	Internal Error. Wait a few minutes, then try the command again
14090343	Invalid Field

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

snaplock_legal_hold_operation

Name	Type	Description
id	integer	Operation ID.
num_files_failed	string	Specifies the number of files on which legal-hold operation failed.
num_files_processed	string	Specifies the number of files on which legal-hold operation was successful.
num_files_skipped	string	Specifies the number of files on which legal-hold begin operation was skipped. The legal-hold begin operation is skipped on a file if it is already under hold for a given litigation.
num_inodes_ignored	string	Specifies the number of inodes on which the legal-hold operation was not attempted because they were not regular files.
path	string	Specifies the path on which legal-hold operation is applied.
state	string	Specifies the status of legal-hold operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Cancel the ongoing legal-hold operation

DELETE /storage/snaplock/litigations/{litigation.id}/operations/{id}

Introduced In: 9.7

Aborts the ongoing legal-hold operation. An abort does not rollback any changes already made. You must re-run begin or end for cleanup.

Related ONTAP commands

- `snaplock legal-hold abort`

Example

```
DELETE "/api/storage/snaplock/litigations/f8a67b60-4461-11e9-b327-0050568ebef5:11/operations/16908292"
```

Learn more

- [DOC /storage/snaplock/litigations/{litigation.id}/operations](#)

Parameters

Name	Type	In	Required	Description
litigation.id	string	path	True	Litigation ID
id	string	path	True	Operation ID.

Response

```
Status: 200, Ok
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090346	Internal Error. Wait a few minutes, then try the command again
14090541	A completed or failed operation cannot be aborted

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the status of legal-hold for an operation ID

GET /storage/snaplock/litigations/{litigation.id}/operations/{id}

Introduced In: 9.7

Retrieves the status of legal-hold for the specified operation ID.

Related ONTAP commands

- `snaplock legal-hold show`

Learn more

- [DOC /storage/snaplock/litigations/{litigation.id}/operations](#)

Parameters

Name	Type	In	Required	Description
litigation.id	string	path	True	Litigation ID
id	string	path	True	Operation ID.

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
id	integer	Operation ID.
num_files_failed	string	Specifies the number of files on which legal-hold operation failed.
num_files_processed	string	Specifies the number of files on which legal-hold operation was successful.
num_files_skipped	string	Specifies the number of files on which legal-hold begin operation was skipped. The legal-hold begin operation is skipped on a file if it is already under hold for a given litigation.
num_inodes_ignored	string	Specifies the number of inodes on which the legal-hold operation was not attempted because they were not regular files.
path	string	Specifies the path on which legal-hold operation is applied.
state	string	Specifies the status of legal-hold operation.
type	string	Specifies the type of legal-hold operation.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "id": "16842759",
  "num_files_failed": "0",
  "num_files_processed": "30",
  "num_files_skipped": "10",
  "num_inodes_ignored": "10",
  "path": "/dir1",
  "state": "completed",
  "type": "begin"
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
14090346	Internal Error. Wait a few minutes, then try the command again
14090343	Invalid Field

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

SnapMirror

SnapMirror overview

Overview

SnapMirror technology enables customers to copy and backup their production data. The secondary or destination volume in such a backup relationship can reside anywhere, locally or remotely, and can be used to restore access to the protected data. You can restore access to protected data by drawing upon a vault of backups or, in the event that a primary is unusable, by using a disaster recovery copy that can be activated for clients and applications. SnapMirror periodically updates a replica to create new backups and/or to keep a replica up-to-date with changes that have been written to the primary. The SnapMirror subsystems are designed to keep many pairs of source (primary) and destination (secondary) copies up-to-date in an efficient and scalable manner.

The SnapMirror APIs can be used to create and manage SnapMirror relationships of type "async", and "sync". These APIs can also be used to manage restore-relationships. These APIs allow you to manage the following endpoints:

- SnapMirror policies - When applied to a SnapMirror relationship, the SnapMirror policy controls the behavior of the relationship and specifies the configuration attributes for that relationship.
- SnapMirror relationships - You can create and manage SnapMirror relationships, and you can change the state of the SnapMirror relationship using a PATCH request.
- SnapMirror transfers - You can manage data transfers on the specified SnapMirror relationship.

Manage SnapMirror policies

SnapMirror policies endpoint overview

Managing SnapMirror policies

This API is used to manage SnapMirror policies of type "async" and "sync". When applied to a SnapMirror relationship, the SnapMirror policy controls the behavior of the relationship and specifies the configuration attributes for that relationship.

The policy type "async" can be associated with a SnapMirror relationship that has either the FlexVol volume or FlexGroup volume or SVM as the endpoint.

The policy type "sync" can be associated with a SnapMirror relationship that has FlexVol volume or a Consistency Group as the endpoint. The policy type "sync" can have a "sync_type" of either "sync", "strict_sync" or "automated_failover". If the "sync_type" is "sync" then a write success is returned to the client after writing the data to the primary endpoint and before writing the data to the secondary endpoint. If the "sync_type" is "strict_sync" then a write success is returned to the client after writing the data to the both primary and secondary endpoints.

The "sync_type" of "automated_failover" can be associated with a SnapMirror relationship that has Consistency Group as the endpoint.

Mapping of SnapMirror policies from CLI to REST

CLI	REST
mirror-vault	async

CLI	REST	sync_type
sync-mirror	sync	sync
strict-sync-mirror	sync	strict_sync
automated-failover	sync	automated_failover

Retrieve SnapMirror async and sync policy types

GET /snapmirror/policies

Introduced In: 9.6

Retrieves SnapMirror policies of type "async" and "sync".

Related ONTAP commands

- `snapmirror policy show`

Example

The following example shows how to retrieve a collection of SnapMirror policies.

```
GET "/api/snapmirror/policies"
```

Learn more

- [DOC /snapmirror/policies](#)

Parameters

Name	Type	In	Required	Description
sync_common_snapshot_schedule.name	string	query	False	Filter by sync_common_snapshot_schedule.name
sync_common_snapshot_schedule.uuid	string	query	False	Filter by sync_common_snapshot_schedule.uuid
type	string	query	False	Filter by type
retention.creation_schedule.name	string	query	False	Filter by retention.creation_schedule.name
retention.creation_schedule.uuid	string	query	False	Filter by retention.creation_schedule.uuid
retention.prefix	string	query	False	Filter by retention.prefix
retention.count	integer	query	False	Filter by retention.count
retention.label	string	query	False	Filter by retention.label
transfer_schedule.name	string	query	False	Filter by transfer_schedule.name

Name	Type	In	Required	Description
transfer_schedule.uuid	string	query	False	Filter by transfer_schedule.uuid
comment	string	query	False	Filter by comment
name	string	query	False	Filter by name
identity_preservation	string	query	False	Filter by identity_preservation
scope	string	query	False	Filter by scope
uuid	string	query	False	Filter by uuid
sync_type	string	query	False	Filter by sync_type
throttle	integer	query	False	Filter by throttle
network_compression_enabled	boolean	query	False	Filter by network_compression_enabled
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[snapmirror_policy]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "identity_preservation": "string",
      "name": "Asynchronous",
      "retention": [
        {
          "count": "7",
          "creation_schedule": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "weekly",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "label": "hourly",
          "prefix": "string"
        }
      ],
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

```

    },
    "sync_common_snapshot_schedule": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "weekly",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "sync_type": "string",
    "throttle": 0,
    "transfer_schedule": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "weekly",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "type": "string",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  }
]
}

```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303842	SnapMirror policy is not supported.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

creation_schedule

Schedule used to create Snapshot copies on the destination for long term retention.

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror_policy_rule

SnapMirror policy rule for retention.

Name	Type	Description
count	integer	Number of Snapshot copies to be kept for retention.
creation_schedule	creation_schedule	Schedule used to create Snapshot copies on the destination for long term retention.
label	string	Snapshot copy label

Name	Type	Description
prefix	string	Specifies the prefix for the Snapshot copy name to be created as per the schedule. If no value is specified, then the label is used as the prefix.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

sync_common_snapshot_schedule

Schedule used to create common Snapshot copies for synchronous relationships.

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

transfer_schedule

The schedule used to update asynchronous relationships.

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror_policy

SnapMirror policy information. SnapMirror policy can either be of type "async" or "sync".

The policy type "async" can be associated with a SnapMirror relationship that has either the FlexVol volume or FlexGroup volume or SVM as the endpoint.

The policy type "sync" along with "sync_type" as "sync" or "strict_sync" can be associated with a SnapMirror relationship that has FlexVol volume as the endpoint. The policy type "sync" can have a "sync_type" of either "sync", "strict_sync" or "automated_failover". If the "sync_type" is "sync" then a write success is returned to the client after writing the data to the source endpoint and before writing the data to

the destination endpoint. If the "sync_type" is "strict_sync" then a write success is returned to the client after writing the data to the both source and destination endpoints.

If the "sync_type" is "automated_failover" then the policy can be associated with a SnapMirror relationship that has Consistency Group as the endpoint. Use the "sync" policy with "sync_type" as "automated_failover" to establish SnapMirror relationships for business continuity usecases. SnapMirror relationships with policy type as "sync" and "sync_type" as "automated_failover" can be monitored by the Mediator, if configured. In case the source Consistency Group endpoint is not reachable, the Mediator may trigger a failover to the destination Consistency Group endpoint.

Name	Type	Description
_links	_links	
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to the destination SVM. This property is applicable only for SVM data protection with "async" policy type.
name	string	
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to the policies of type "async".
retention	array[snapmirror_policy_rule]	Policy on Snapshot copy retention.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	svm	
sync_common_snapshot_schedule	sync_common_snapshot_schedule	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	transfer_schedule	The schedule used to update asynchronous relationships.
type	string	

Name	Type	Description
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a SnapMirror policy

POST /snapmirror/policies

Introduced In: 9.6

Creates a SnapMirror policy. The property "identity_preservation" is applicable to only SnapMirror relationships with SVM endpoints and it indicates which configuration of the source SVM is replicated to the destination SVM.

It takes the following values:

- `full` - indicates that the source SVM configuration is replicated to the destination SVM endpoint.
- `exclude_network_config` - indicates that the source SVM configuration other than network configuration is replicated to the destination SVM endpoint.
- `exclude_network_and_protocol_config` - indicates that the source SVM configuration is not replicated to the destination SVM endpoint.

Important note

- The property "identity_preservation" is applicable to only SnapMirror relationships with SVM endpoints and it indicates which configuration of the source SVM is replicated to the destination SVM.
- The properties "identity_preservation" and "transfer_schedule" are not applicable for "sync" type policies.

- The properties "retention.creation_schedule" and "retention.prefix" are not applicable for "sync" type policies.
- The property "sync_common_snapshot_schedule" is not applicable for an "async" type policy.
- The property "retention.count" specifies the maximum number of Snapshot copies that are retained on the SnapMirror destination volume.
- When the property "retention.label" is specified, the Snapshot copies that have a SnapMirror label matching this property is transferred to the SnapMirror destination.
- When the property "retention.creation_schedule" is specified, Snapshot copies are directly created on the SnapMirror destination. The Snapshot copies created have the same content as the latest Snapshot copy already present on the SnapMirror destination.

Required properties

- `name` - Name of the new SnapMirror policy.

Recommended optional properties

- `svm.name` or `svm.uuid` - Name or UUID of the SVM that owns the SnapMirror policy.

Default property values

If not specified in POST, the following default property values are assigned:

- `type` - *async*
- `sync_type` - *sync* (when `type` is *sync*)
- `network_compression_enabled` - *false*
- `throttle` - *0*
- `identity_preservation` - *exclude_network_and_protocol_config*

Related ONTAP commands

- `snapmirror policy create`

Examples

Creating a SnapMirror policy of type "sync"

```
POST "/api/snapmirror/policies/" '{"name": "policy1", "svm.name": "VS0",
"type": "sync", "sync_type": "sync"}'
```

Creating a SnapMirror policy of type "async" with retention values

```
POST "/api/snapmirror/policies" '{"name": "policy_ret", "svm": {"name":
"vs1"}, "retention": {"label": ["smcreate"], "count": ["2"],
"creation_schedule": ["weekly"]}]'
```

Creating a SnapMirror policy of type "async"

```
POST "/api/snapmirror/policies" '{"name": "newPolicy", "svm":{"name" : "vs1"}, "type": "async"}'
```

Creating a SnapMirror policy of type "sync" with sync_type as "automated_failover"

```
POST "/api/snapmirror/policies/" '{"name": "policy1", "svm.name": "VS0", "type": "sync", "sync_type": "automated_failover" }'
```

Learn more

- [DOC /snapmirror/policies](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<div>The default is false. If set to true, the records are returned.</div> <div>• Default value:</div>

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to the destination SVM. This property is applicable only for SVM data protection with "async" policy type.
name	string	

Name	Type	Description
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to the policies of type "async".
retention	array[snapmirror_policy_rule]	Policy on Snapshot copy retention.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	svm	
sync_common_snapshot_schedule	sync_common_snapshot_schedule	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	transfer_schedule	The schedule used to update asynchronous relationships.
type	string	
uuid	string	

Example request

```
{
  "comment": "string",
  "identity_preservation": "string",
  "name": "Asynchronous",
  "retention": [
    {
      "count": "7",
      "creation_schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "label": "hourly",
      "prefix": "string"
    }
  ],
  "scope": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "sync_common_snapshot_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "sync_type": "string",
  "throttle": 0,
  "transfer_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303850	Invalid input parameter
13303887	Failed to create SnapMirror policy. Reason: Maximum number of allowed retention rules reached

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

creation_schedule

Schedule used to create Snapshot copies on the destination for long term retention.

Name	Type	Description
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror_policy_rule

SnapMirror policy rule for retention.

Name	Type	Description
count	integer	Number of Snapshot copies to be kept for retention.
creation_schedule	creation_schedule	Schedule used to create Snapshot copies on the destination for long term retention.
label	string	Snapshot copy label
prefix	string	Specifies the prefix for the Snapshot copy name to be created as per the schedule. If no value is specified, then the label is used as the prefix.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

sync_common_snapshot_schedule

Schedule used to create common Snapshot copies for synchronous relationships.

Name	Type	Description
name	string	Job schedule name
uuid	string	Job schedule UUID

transfer_schedule

The schedule used to update asynchronous relationships.

Name	Type	Description
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror_policy

SnapMirror policy information. SnapMirror policy can either be of type "async" or "sync".

The policy type "async" can be associated with a SnapMirror relationship that has either the FlexVol volume or FlexGroup volume or SVM as the endpoint.

The policy type "sync" along with "sync_type" as "sync" or "strict_sync" can be associated with a SnapMirror relationship that has FlexVol volume as the endpoint. The policy type "sync" can have a "sync_type" of either "sync", "strict_sync" or "automated_failover". If the "sync_type" is "sync" then a write success is returned to the client after writing the data to the source endpoint and before writing the data to the destination endpoint. If the "sync_type" is "strict_sync" then a write success is returned to the client after writing the data to the both source and destination endpoints.

If the "sync_type" is "automated_failover" then the policy can be associated with a SnapMirror relationship that has Consistency Group as the endpoint. Use the "sync" policy with "sync_type" as "automated_failover" to establish SnapMirror relationships for business continuity usecases. SnapMirror relationships with policy type as "sync" and "sync_type" as "automated_failover" can be monitored by the Mediator, if configured. In case the source Consistency Group endpoint is not reachable, the Mediator may trigger a failover to the destination Consistency Group endpoint.

Name	Type	Description
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to the destination SVM. This property is applicable only for SVM data protection with "async" policy type.
name	string	

Name	Type	Description
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to the policies of type "async".
retention	array[snapmirror_policy_rule]	Policy on Snapshot copy retention.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	svm	
sync_common_snapshot_schedule	sync_common_snapshot_schedule	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	transfer_schedule	The schedule used to update asynchronous relationships.
type	string	
uuid	string	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a SnapMirror policy

DELETE /snapmirror/policies/{uuid}

Introduced In: 9.6

Deletes a SnapMirror policy.

Related ONTAP commands

- `snapmirror policy delete`

Example

```
DELETE "/api/snapmirror/policies/510c15d4-f9e6-11e8-bdb5-0050568e12c2"
```

Learn more

- [DOC /snapmirror/policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Policy UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a specific SnapMirror policy

GET /snapmirror/policies/{uuid}

Introduced In: 9.6

Retrieves a specific SnapMirror policy.

Example

```
GET "/api/snapmirror/policies/567aaac0-f863-11e8-a666-0050568e12c2"
```

Learn more

- [DOC /snapmirror/policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Policy UUID
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to the destination SVM. This property is applicable only for SVM data protection with "async" policy type.
name	string	
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to the policies of type "async".
retention	array[snapmirror_policy_rule]	Policy on Snapshot copy retention.

Name	Type	Description
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
svm	svm	
sync_common_snapshot_schedule	sync_common_snapshot_schedule	Schedule used to create common Snapshot copies for synchronous relationships.
sync_type	string	
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	transfer_schedule	The schedule used to update asynchronous relationships.
type	string	
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "identity_preservation": "string",
  "name": "Asynchronous",
  "retention": [
    {
      "count": "7",
      "creation_schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "label": "hourly",
      "prefix": "string"
    }
  ],
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "sync_common_snapshot_schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
}
```

```
"sync_type": "string",
"throttle": 0,
"transfer_schedule": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "weekly",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"type": "string",
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303842	SnapMirror policy is not supported.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

creation_schedule

Schedule used to create Snapshot copies on the destination for long term retention.

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror_policy_rule

SnapMirror policy rule for retention.

Name	Type	Description
count	integer	Number of Snapshot copies to be kept for retention.
creation_schedule	creation_schedule	Schedule used to create Snapshot copies on the destination for long term retention.
label	string	Snapshot copy label
prefix	string	Specifies the prefix for the Snapshot copy name to be created as per the schedule. If no value is specified, then the label is used as the prefix.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

sync_common_snapshot_schedule

Schedule used to create common Snapshot copies for synchronous relationships.

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

transfer_schedule

The schedule used to update asynchronous relationships.

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Update the SnapMirror policy

PATCH /snapmirror/policies/{uuid}

Introduced In: 9.6

Updates the SnapMirror policy.

Important notes

- The properties "transfer_schedule" and "throttle" can be modified only if all the SnapMirror relationships associated with the specified SnapMirror policy have the same values.
- The properties "retention.label" and "retention.count" are mandatory if "retention" is provided in the input. The provided "retention.label" is the final list and is replaced with the existing values.
- The value of the "identity_preservation" property cannot be changed if the SnapMirror relationships associated with the policy have different identity_preservation configurations.
- If the SnapMirror policy "identity_preservation" value matches the "identity_preservation" value of the associated SnapMirror relationships, then the "identity_preservation" value can be changed from a higher "identity_preservation" threshold value to a lower "identity_preservation" threshold value but not vice-versa. For example, the threshold value of the "identity_preservation" property can be changed from "full" to "exclude_network_config" to "exclude_network_and_protocol_config", but could not be increased from "exclude_network_and_protocol_config" to "exclude_network_config" to "full".

Related ONTAP commands

- `snapmirror policy modify`

Example

Updating the "retention" property

```
PATCH "/api/snapmirror/policies/fe65686d-00dc-11e9-b5fb-0050568e3f83"
'{"retention" : {"label" : ["sm_created", "lab2"], "count": ["1","2"],
"creation_schedule": {"name": ["weekly"]}}}'
```

Updating "transfer_schedule", "throttle", and "identity_preservation" properties

```
PATCH "/api/snapmirror/policies/8aef950b-3bef-11e9-80ac-0050568ea591"
'{"transfer_schedule.name" : "weekly", "throttle" : "100",
"identity_preservation":"exclude_network_and_protocol_config"}'
```

Learn more

- [DOC /snapmirror/policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Policy UUID
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Request Body

Name	Type	Description
comment	string	Comment associated with the policy.

Name	Type	Description
identity_preservation	string	Specifies which configuration of the source SVM is replicated to the destination SVM. This property is applicable only for SVM data protection with "async" policy type.
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to the policies of type "async".
retention	array[snapmirror_policy_rule]	Policy on Snapshot copy retention.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
sync_common_snapshot_schedule	sync_common_snapshot_schedule	Schedule used to create common Snapshot copies for synchronous relationships.
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	transfer_schedule	The schedule used to update asynchronous relationships.
uuid	string	

Example request

```
{
  "comment": "string",
  "identity_preservation": "string",
  "retention": [
    {
      "count": "7",
      "creation_schedule": {
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "label": "hourly",
      "prefix": "string"
    }
  ],
  "scope": "string",
  "sync_common_snapshot_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "throttle": 0,
  "transfer_schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303842	SnapMirror policy is not supported.
13303843	Conflicting values between SnapMirror policy and SnapMirror relationships for either 'transfer_schedule, throttle or identity_preservation' properties
13303850	Invalid input parameter
13303887	Failed to create SnapMirror policy. Reason: Maximum number of allowed retention rules reached

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

creation_schedule

Schedule used to create Snapshot copies on the destination for long term retention.

Name	Type	Description
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror_policy_rule

SnapMirror policy rule for retention.

Name	Type	Description
count	integer	Number of Snapshot copies to be kept for retention.
creation_schedule	creation_schedule	Schedule used to create Snapshot copies on the destination for long term retention.
label	string	Snapshot copy label
prefix	string	Specifies the prefix for the Snapshot copy name to be created as per the schedule. If no value is specified, then the label is used as the prefix.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

sync_common_snapshot_schedule

Schedule used to create common Snapshot copies for synchronous relationships.

Name	Type	Description
name	string	Job schedule name
uuid	string	Job schedule UUID

transfer_schedule

The schedule used to update asynchronous relationships.

Name	Type	Description
name	string	Job schedule name
uuid	string	Job schedule UUID

snapmirror_policy

SnapMirror policy information. SnapMirror policy can either be of type "async" or "sync".

The policy type "async" can be associated with a SnapMirror relationship that has either the FlexVol volume or FlexGroup volume or SVM as the endpoint.

The policy type "sync" along with "sync_type" as "sync" or "strict_sync" can be associated with a SnapMirror relationship that has FlexVol volume as the endpoint. The policy type "sync" can have a "sync_type" of either "sync", "strict_sync" or "automated_failover". If the "sync_type" is "sync" then a write success is returned to the client after writing the data to the source endpoint and before writing the data to the destination endpoint. If the "sync_type" is "strict_sync" then a write success is returned to the client after writing the data to the both source and destination endpoints.

If the "sync_type" is "automated_failover" then the policy can be associated with a SnapMirror relationship that has Consistency Group as the endpoint. Use the "sync" policy with "sync_type" as "automated_failover" to establish SnapMirror relationships for business continuity usecases. SnapMirror relationships with policy type as "sync" and "sync_type" as "automated_failover" can be monitored by the Mediator, if configured. In case the source Consistency Group endpoint is not reachable, the Mediator may trigger a failover to the destination Consistency Group endpoint.

Name	Type	Description
comment	string	Comment associated with the policy.
identity_preservation	string	Specifies which configuration of the source SVM is replicated to the destination SVM. This property is applicable only for SVM data protection with "async" policy type.

Name	Type	Description
network_compression_enabled	boolean	Specifies whether network compression is enabled for transfers. This is applicable only to the policies of type "async".
retention	array[snapmirror_policy_rule]	Policy on Snapshot copy retention.
scope	string	Set to "svm" for policies owned by an SVM, otherwise set to "cluster".
sync_common_snapshot_schedule	sync_common_snapshot_schedule	Schedule used to create common Snapshot copies for synchronous relationships.
throttle	integer	Throttle in KB/s. Default to unlimited.
transfer_schedule	transfer_schedule	The schedule used to update asynchronous relationships.
uuid	string	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage SnapMirror relationships

SnapMirror relationships endpoint overview

Overview

This API manages asynchronous extended data protection (XDP) relationships for FlexVols, FlexGroups, or SVMs. It is also used to manage a synchronous relationship between FlexVol volumes, which provides zero RPO data protection and active synchronous relationship with automated failover between Consistency Group endpoints which provides zero RTO data protection. To create an asynchronous extended data protection relationship with FlexVol volumes, FlexGroup volumes, or SVMs, use the policy of type "async". To create a synchronous relationship between FlexVol volumes, use the policy of type "sync" with sync_type of either "sync" or "strict_sync". To create an active synchronous relationship with automated failover between Consistency Group endpoints, use the policy of type "sync" with sync_type "automated_failover". You can create an asynchronous extended data protection relationship between the source and destination which can be used by the transfer APIs to perform SnapMirror "restore" operations.

To create FlexVol or FlexGroup SnapMirror relationships, the source volume must be in the "online" state and be a read-write type; the destination volume must be in the "online" state and be a data protection type. To create SnapMirror relationships between SVMs, the source SVM must be of subtype "default" and the destination SVM of subtype "dp_destination". Additionally, SVMs must be peered before a relationship can be established between them when the "create_destination" property is not specified. When the "create_destination" property is specified then the destination SVM is provisioned on the destination cluster and the SVM peer relationship is established between the source SVM and the new destination SVM provided the source SVM has the SVM peering permission for the destination cluster. The SnapMirror functionality is subdivided into relationship APIs and transfer APIs:

- SnapMirror relationship APIs are used to create and manage the SnapMirror relationships.
- SnapMirror transfer APIs are used to manage data transfers.

Retrieve information for SnapMirror relationships

GET /snapmirror/relationships

Introduced In: 9.6

Retrieves information for SnapMirror relationships whose destination endpoints are in the current SVM or the current cluster, depending on the cluster context.

Related ONTAP commands

- `snapmirror show`

- `snapmirror list-destinations`

Examples

The following examples show how to retrieve the list of SnapMirror relationships and the list of SnapMirror destinations.

1. Retrieving the list of SnapMirror relationships. This API must be run on the cluster containing the destination endpoint.

```
GET "/api/snapmirror/relationships/"
```

1. Retrieving the list of SnapMirror destinations on source. This must be run on the cluster containing the source endpoint.

```
GET "/api/snapmirror/relationships/?list_destinations_only=true"
```

Learn more

- [DOC /snapmirror/relationships](#)

Parameters

Name	Type	In	Required	Description
list_destinations_only	boolean	query	False	Set to true to show relationships from the source only.
transfer.state	string	query	False	Filter by transfer.state
transfer.bytes_transferred	integer	query	False	Filter by transfer.bytes_transferred
transfer.uuid	string	query	False	Filter by transfer.uuid
destination.cluster.name	string	query	False	Filter by destination.cluster.name <ul style="list-style-type: none"> • Introduced in: 9.7

Name	Type	In	Required	Description
destination.cluster.uuid	string	query	False	Filter by destination.cluster.uuid • Introduced in: 9.7
destination.consistency_group_volumes.uuid	string	query	False	Filter by destination.consistency_group_volumes.uuid • Introduced in: 9.8
destination.consistency_group_volumes.name	string	query	False	Filter by destination.consistency_group_volumes.name • Introduced in: 9.8
destination.path	string	query	False	Filter by destination.path
destination.svm.uuid	string	query	False	Filter by destination.svm.uuid
destination.svm.name	string	query	False	Filter by destination.svm.name
uuid	string	query	False	Filter by uuid
consistency_group_failover.error.code	string	query	False	Filter by consistency_group_failover.error.code • Introduced in: 9.8
consistency_group_failover.error.arguments.code	string	query	False	Filter by consistency_group_failover.error.arguments.code • Introduced in: 9.8

Name	Type	In	Required	Description
consistency_group_failover.error.arguments.message	string	query	False	Filter by consistency_group_failover.error.arguments.message <ul style="list-style-type: none"> • Introduced in: 9.8
consistency_group_failover.error.message	string	query	False	Filter by consistency_group_failover.error.message <ul style="list-style-type: none"> • Introduced in: 9.8
consistency_group_failover.error.target	string	query	False	Filter by consistency_group_failover.error.target <ul style="list-style-type: none"> • Introduced in: 9.8
consistency_group_failover.status.message	string	query	False	Filter by consistency_group_failover.status.message <ul style="list-style-type: none"> • Introduced in: 9.8
consistency_group_failover.status.code	string	query	False	Filter by consistency_group_failover.status.code <ul style="list-style-type: none"> • Introduced in: 9.8
source.cluster.name	string	query	False	Filter by source.cluster.name <ul style="list-style-type: none"> • Introduced in: 9.7
source.cluster.uuid	string	query	False	Filter by source.cluster.uuid <ul style="list-style-type: none"> • Introduced in: 9.7

Name	Type	In	Required	Description
source.consistency_group_volumes.uuid	string	query	False	Filter by source.consistency_group_volumes.uuid • Introduced in: 9.8
source.consistency_group_volumes.name	string	query	False	Filter by source.consistency_group_volumes.name • Introduced in: 9.8
source.path	string	query	False	Filter by source.path
source.svm.uuid	string	query	False	Filter by source.svm.uuid
source.svm.name	string	query	False	Filter by source.svm.name
lag_time	string	query	False	Filter by lag_time
policy.type	string	query	False	Filter by policy.type
policy.uuid	string	query	False	Filter by policy.uuid
policy.name	string	query	False	Filter by policy.name
state	string	query	False	Filter by state
unhealthy_reason.parameters	string	query	False	Filter by unhealthy_reason.parameters
unhealthy_reason.message	string	query	False	Filter by unhealthy_reason.message
unhealthy_reason.code	integer	query	False	Filter by unhealthy_reason.code
restore	boolean	query	False	Filter by restore

Name	Type	In	Required	Description
exported_snapshot	string	query	False	Filter by exported_snapshot
healthy	boolean	query	False	Filter by healthy
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[snapmirror_relationship]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "consistency_group_failover": {
        "error": {
          "arguments": [
            {
              "code": "string",
              "message": "string"
            }
          ],
          "code": "4",
          "message": "entry doesn't exist",
          "target": "uuid"
        },
        "status": {
          "code": "string",
          "message": "string"
        }
      },
      "destination": {
        "cluster": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "consistency_group_volumes": [
```

```

    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ],
  "path": "svm1:volume1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
},
"exported_snapshot": "string",
"lag_time": "PT8H35M42S",
"policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "Asynchronous",
  "type": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
},
"restore_to_snapshot": "string",
"source": {
  "cluster": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "cluster1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "consistency_group_volumes": [
    {

```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "path": "svm1:volume1",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "state": "snapmirrored",
  "transfer": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "state": "string",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "unhealthy_reason": [
    {
      "code": "6621444",
      "message": "Failed to complete update operation on one or
more item relationships.",
      "parameters": []
    },
    {
      "code": "6621445",
      "message": "Group Update failed",
      "parameters": []
    }
  ],
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
]

```

```
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303825	Could not retrieve information for the SnapMirror policy type
13303817	Unknown value for the Snapmirror State

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

status

Name	Type	Description
code	string	Status code

Name	Type	Description
message	string	SnapMirror Consistency Group failover status.

snapmirror_consistency_group_failover

SnapMirror Consistency Group failover information. The SnapMirror Consistency Group failover can be a planned or an unplanned operation. Only active SnapMirror Consistency Group failover operation progress can be monitored using this object. In case of an error during the failover operation, the property "consistency_group_failover.error" holds the reason for the error. ONTAP automatically retries any failed SnapMirror Consistency Group failover operation.

Name	Type	Description
error	error	
status	status	

storage_service

Name	Type	Description
enabled	boolean	This property indicates whether to create the destination endpoint using storage service.
enforce_performance	boolean	Optional property to enforce storage service performance on the destination endpoint. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints.

Name	Type	Description
name	string	<p>Optional property to specify the storage service name for the destination endpoint. This property is considered when the property "create_destination.storage_service.enabled" is set to "true". When the property "create_destination.storage_service.enabled" is set to "true" and the "create_destination.storage_service.name" for the endpoint is not specified, then ONTAP selects the highest storage service available on the cluster to provision the destination endpoint. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints.</p> <ul style="list-style-type: none"> enum: ["extreme", "performance", "value"] Introduced in: 9.6

tiering

Name	Type	Description
policy	string	<p>Optional property to specify the destination endpoint's tiering policy when "create_destination.tiering.supported" is set to "true". This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. This property determines whether the user data blocks of the destination endpoint in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of the destination endpoint volume blocks increases if they are accessed frequently and decreases when they are not.</p> <p>all – This policy allows tiering of both destination endpoint Snapshot copies and the user transferred data blocks to the cloud store as soon as possible by ignoring the temperature on the volume blocks. This tiering policy is not applicable for Consistency Group destination endpoints or for synchronous relationships.</p> <p>auto – This policy allows tiering of both destination endpoint Snapshot copies and the active file system user data to the cloud store</p> <p>none – Destination endpoint volume blocks will not be tiered to the cloud store.</p> <p>snapshot_only – This policy allows tiering of only the destination endpoint volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot_only" for a FlexVol volume and "none" for a FlexGroup volume.</p>

Name	Type	Description
supported	boolean	Optional property to enable provisioning of the destination endpoint volumes on FabricPool aggregates. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. Only FabricPool aggregates are used if this property is set to "true" and only non FabricPool aggregates are used if this property is set to "false". Tiering support for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only do the destination endpoint volumes need to support tiering by using FabricPools, the "create_destination.tiering.policy" must not be "none". A destination endpoint that uses FabricPools but has a tiering "policy" of "none" supports tiering but will not tier any data.

snapmirror_destination_creation

Use this object to provision the destination endpoint when establishing a SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM, or Consistency Group. Given a source endpoint, the destination endpoint is provisioned in the SVM specified in the "destination.path" property. While protecting an SVM, the SVM destination endpoint can only be provisioned on the local cluster. To provision the SVM destination endpoint use the optional "source.cluster.name" property to specify the remote cluster name or use the optional "source.cluster.uuid" property to specify the remote cluster UUID. When "create_destination.enabled" option is specified while making a POST for a SnapMirror relationship, the relationship can be automatically initialized by setting the "state" either to "snapmirrored" when the policy is of type "async" or to "in_sync" when the policy is of type "sync". The "destination.path" property must specify the destination endpoint path. For example, for FlexVol volume and FlexGroup volume, the "destination.path" can be specified as <destination-svm-name:dp-volume-name>, for SVM data protection, the "destination.path" must be specified as <destination-svm-name:>, and for Consistency Group, the "destination.path" must be specified as <destination-svm-name:> along with the "destination.consistency_group_volumes" property to indicate the list of destination volumes of type "DP" in the Consistency Group. For a FlexVol volume, a FlexGroup volume, or a Consistency Group destination endpoint, the properties in this object can be specified either from the source or the destination cluster. For an SVM destination endpoint, the properties in this object can be specified from the destination cluster only. This object is not supported for non ONTAP endpoints.</destination-svm-name:></destination-svm-name:></destination-svm-name:dp-volume-name>

Name	Type	Description
enabled	boolean	Optional property to create the destination endpoint when establishing a SnapMirror relationship. It is assumed to be "false" if no other property is set and assumed to be "true" if any other property is set.
storage_service	storage_service	
tiering	tiering	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

consistency_group_volumes

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the

endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
cluster	cluster	
consistency_group_volumes	array[consistency_group_volumes]	Mandatory property for a Consistency Group endpoint. Specifies the list of FlexVol volumes for a Consistency Group.
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: ONTAP Consistency Group - svm1:/cg/cg_name <ul style="list-style-type: none">• example: svm1:volume1• Introduced in: 9.6
svm	svm	

policy

Basic policy information of the relationship.

Name	Type	Description
_links	_links	
name	string	
type	string	
uuid	string	

transfer

Basic information on the current transfer.

Name	Type	Description
_links	_links	
bytes_transferred	integer	Bytes transferred.
state	string	

Name	Type	Description
uuid	string	

snapmirror_error

SnapMirror error

Name	Type	Description
code	integer	Error code
message	string	Error message
parameters	array[string]	Parameters for the error message

snapmirror_relationship

SnapMirror relationship information. The SnapMirror relationship can be either "async" or "sync" based on the type of SnapMirror policy associated with the relationship. The source and destination endpoints of a SnapMirror relationship must be of the same type, for example, if the source endpoint is a FlexVol volume then the destination endpoint must be a FlexVol volume.

The SnapMirror policy type "async" can be used when the SnapMirror relationship has FlexVol volume or FlexGroup volume or SVM as the endpoint. The SnapMirror policy type "sync" can be used when the SnapMirror relationship has FlexVol volume as the endpoint. The SnapMirror policy type "sync" with "sync_type" as "automated_failover" can be used when the SnapMirror relationship has Consistency Group as the endpoint.

Name	Type	Description
_links	_links	
consistency_group_failover	snapmirror_consistency_group_failover	<p>SnapMirror Consistency Group failover information. The SnapMirror Consistency Group failover can be a planned or an unplanned operation. Only active SnapMirror Consistency Group failover operation progress can be monitored using this object. In case of an error during the failover operation, the property "consistency_group_failover.error" holds the reason for the error. ONTAP automatically retries any failed SnapMirror Consistency Group failover operation.</p> <ul style="list-style-type: none"> Introduced in: 9.8 readOnly: 1

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
exported_snapshot	string	Snapshot copy exported to clients on destination.
healthy	boolean	Is the relationship healthy?
lag_time	string	Time since the exported Snapshot copy was created.
policy	policy	Basic policy information of the relationship.
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This property is applicable only for relationships with FlexVol volume or FlexGroup volume endpoints and when the PATCH state is being changed to "snapmirrored".

Name	Type	Description
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This property is applicable only for relationships with FlexVol volume endpoints and when the PATCH state is being changed to "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup volume relationship. This restores all destination FlexGroup constituent volumes to the latest Snapshot copy, and any writes to the read-write constituents are lost. This property is applicable only for SnapMirror relationships with FlexGroup volume endpoints and when the PATCH state is being changed to "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship. SnapMirror relationships with the policy type "async" can be restored. SnapMirror relationships with the policy type "sync" cannot be restored.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination during the break operation. This property is applicable only for SnapMirror relationships with FlexVol volume endpoints and when the PATCH state is being changed to "broken_off".

Name	Type	Description
source	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
state	string	<p>State of the relationship.</p> <p>To initialize the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>To break the relationship, PATCH the state to "broken_off" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "broken_off".</p> <p>To resync the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" can be in "broken_off" state due to a failed attempt of SnapMirror failover.</p> <p>To pause the relationship, suspending further transfers, PATCH the state to "paused" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "paused".</p> <p>To resume transfers for a paused relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships with a policy of type "sync". A PATCH call on the state change only triggers the transition to the specified state. You must poll on the "state", "healthy" and "unhealthy_reason" properties using a GET request to determine if the transition is successful. To automatically initialize the relationship when specifying</p>

Name	Type	Description
transfer	transfer	Basic information on the current transfer.
unhealthy_reason	array[snapmirror_error]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

Create a SnapMirror relationship

POST /snapmirror/relationships

Introduced In: 9.6

Creates a SnapMirror relationship. This API can optionally provision the destination endpoint when it does not exist. This API must be executed on the cluster containing the destination endpoint unless the destination endpoint is being provisioned. When the destination endpoint is being provisioned, this API can also be executed from the cluster containing the source endpoint. Provisioning of the destination endpoint from the source cluster is supported for the FlexVol volume, FlexGroup volume and Consistency Group endpoints. For SVM endpoint, provisioning of the destination SVM endpoint is not supported from the source cluster. When the destination endpoint exists the source SVM and the destination SVM must be in an SVM peer relationship. When provisioning the destination endpoint, the SVM peer relationship between the source SVM and the destination SVM is established as part of the destination, provision provided the source SVM has SVM peering permission for the destination cluster.

Required properties

- `source.path` - Path to the source endpoint of the SnapMirror relationship.
- `destination.path` - Path to the destination endpoint of the SnapMirror relationship.
- `destination.consistency_group_volumes` - List of FlexVol volumes of type "RW" that are constituents of a Consistency Group.
- `destination.consistency_group_volumes` - List of FlexVol volumes of type "DP" that are constituents of a Consistency Group.

Recommended optional properties

- `policy.name` or `policy.uuid` - Policy governing the SnapMirror relationship.
- `state` - Set the state to "snapmirrored" to automatically initialize the relationship.
- `create_destination.enabled` - Enable this property to provision the destination endpoint.

Default property values

If not specified in POST, the following default property values are assigned:

- `policy.name` - *Asynchronous*

- `restore` - *false*
- `create_destination.tiering.policy` - *snapshot_only* (when `create_destination.tiering.supported` is *true* for FlexVol volume)
- `create_destination.tiering.policy` - *none* (when `create_destination.tiering.supported` is *true* for FlexGroup volume)
- `create_destination.storage_service.enforce_performance` - *false*
- `source.ipspace` - *Default*
- `destination.ipspace` - *Default*

Related ONTAP commands

- `snapmirror create`
- `snapmirror protect`

Examples

The following examples show how to create FlexVol, FlexGroup, SVM and Consistency Group SnapMirror relationships. Note that the source SVM name should be the local name of the peer SVM.

Creating a FlexVol SnapMirror relationship of type XDP.

```
POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"}}'
```

Creating a FlexGroup SnapMirror relationship of type XDP.

```
POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:source_flexgrp"}, "destination": { "path":
"dst_svm:dest_flexgrp"}}'
```

Creating a SVM SnapMirror relationship of type XDP.

```
POST "/api/snapmirror/relationships/" '{"source": { "path": "src_svm:"},
"destination": { "path": "dst_svm:"}}'
```

Creating a SnapMirror relationship in order to restore from a destination.

```
POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"restore": "true"}'
```

Provision the destination FlexVol volume endpoint and create a SnapMirror relationship of type XDP.

```
POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"create_destination": { "enable": "true" }}'
```

Provision the destination FlexVol volume endpoint on a Fabricpool with a tiering policy and create a SnapMirror relationship of type XDP.

```
POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"create_destination": { "enable": "true", "tiering": { "supported":
"true", "policy": "auto" } } }'
```

Provision the destination FlexVol volume endpoint using storage service and create a SnapMirror relationship of type XDP.

```
POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:src_vol"}, "destination": { "path": "dst_svm:dst_vol"},
"create_destination": { "enable": "true", "storage_service": { "enabled":
"true", "name": "extreme", "enforce_performance": "true" } } }'
```

Provision the destination SVM endpoint and create a SnapMirror relationship of type XDP.

```
POST "/api/snapmirror/relationships/" '{"source": {"path": "src_svm:",
"cluster": { "name": "cluster_src" }}, "destination": { "path":
"dst_svm:", "create_destination": { "enable": "true" }}'
```

Create a SnapMirror relationship with Consistency Group endpoint.

```
POST "/api/snapmirror/relationships/" '{"source": { "path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": "src_vol_1,
src_vol_2"}, "destination": { "path": "dst_svm:/cg/cg_dst_vol",
"consistency_group_volumes": "dst_vol_1, dst_vol_2"}, "policy":
"AutomatedFailOver" }'
```

Provision the destination Consistency Group endpoint on a Fabricpool with a tiering policy, create a SnapMirror relationship with a SnapMirror policy of type "sync" and sync_type of "automated_failover", and initialize the SnapMirror relationship with state as "in_sync".

```
POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": "src_vol_1,
src_vol_2"}, "destination": { "path": "dst_svm:/cg/cg_dst_vol",
"consistency_group_volumes": "dst_vol_1, dst_vol_2"},
"create_destination": { "enable": "true", "tiering": { "supported": "true"
} }, "policy": "AutomatedFailOver", "state": "in_sync" }'
```

Provision the destination Consistency Group endpoint with storage service, create a SnapMirror relationship with a SnapMirror policy of type "sync" and sync_type of "automated_failover", and initialize the SnapMirror relationship with state as "in_sync".

```
POST "/api/snapmirror/relationships/" '{"source": {"path":
"src_svm:/cg/cg_src_vol", "consistency_group_volumes": "src_vol_1,
src_vol_2"}, "destination": { "path": "dst_svm:/cg/cg_dst_vol",
"consistency_group_volumes": "dst_vol_1, dst_vol_2"},
"create_destination": { "enable": "true", "storage_service": { "enabled":
"true", "name": "extreme", "enforce_performance": "true" } }, "policy":
"AutomatedFailOver", "state": "in_sync" }'
```

Learn more

- [DOC /snapmirror/relationships](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
validate_only	boolean	query	False	<p>Validate the operation and its parameters, without actually performing the operation.</p> <ul style="list-style-type: none"> • Introduced in: 9.7

Request Body

Name	Type	Description
create_destination	snapmirror_destination_creation	<p>Use this object to provision the destination endpoint when establishing a SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM, or Consistency Group. Given a source endpoint, the destination endpoint is provisioned in the SVM specified in the "destination.path" property. While protecting an SVM, the SVM destination endpoint can only be provisioned on the local cluster. To provision the SVM destination endpoint use the optional "source.cluster.name" property to specify the remote cluster name or use the optional "source.cluster.uuid" property to specify the remote cluster UUID. When "create_destination.enabled" option is specified while making a POST for a SnapMirror relationship, the relationship can be automatically initialized by setting the "state" either to "snapmirrored" when the policy is of type "async" or to "in_sync" when the policy is of type "sync". The "destination.path" property must specify the destination endpoint path. For example, for FlexVol volume and FlexGroup volume, the "destination.path" can be specified as <destination-SVM-name:dp-volume-name>, for SVM data protection, the "destination.path" must be specified as &lt;destination-SVM-name:&gt;, and for Consistency Group, the "destination.path" must be specified as &lt;destination-SVM-name:/cg/consistency-group-name&gt; along with the "destination.consistency_group_volumes" property to indicate the list of destination volumes of type "DP" in the Consistency Group. For a FlexVol volume, a FlexGroup volume, or a Consistency Group destination endpoint, the properties in this object can be specified either from the source or the destination cluster. For an SVM destination endpoint, the properties in this object can be specified from the destination cluster only. This</p>

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
exported_snapshot	string	Snapshot copy exported to clients on destination.
healthy	boolean	Is the relationship healthy?
lag_time	string	Time since the exported Snapshot copy was created.
policy	policy	Basic policy information of the relationship.
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This property is applicable only for relationships with FlexVol volume or FlexGroup volume endpoints and when the PATCH state is being changed to "snapmirrored".

Name	Type	Description
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This property is applicable only for relationships with FlexVol volume endpoints and when the PATCH state is being changed to "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup volume relationship. This restores all destination FlexGroup constituent volumes to the latest Snapshot copy, and any writes to the read-write constituents are lost. This property is applicable only for SnapMirror relationships with FlexGroup volume endpoints and when the PATCH state is being changed to "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship. SnapMirror relationships with the policy type "async" can be restored. SnapMirror relationships with the policy type "sync" cannot be restored.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination during the break operation. This property is applicable only for SnapMirror relationships with FlexVol volume endpoints and when the PATCH state is being changed to "broken_off".

Name	Type	Description
source	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
state	string	<p>State of the relationship.</p> <p>To initialize the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>To break the relationship, PATCH the state to "broken_off" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "broken_off".</p> <p>To resync the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" can be in "broken_off" state due to a failed attempt of SnapMirror failover.</p> <p>To pause the relationship, suspending further transfers, PATCH the state to "paused" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "paused".</p> <p>To resume transfers for a paused relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships with a policy of type "sync". A PATCH call on the state change only triggers the transition to the specified state. You must poll on the "state", "healthy" and "unhealthy_reason" properties using a GET request to determine if the transition is successful. To automatically initialize the relationship when specifying "create_destination" property, set the state to "snapmirrored" for</p>

Name	Type	Description
transfer	transfer	Basic information on the current transfer.
unhealthy_reason	array[snapmirror_error]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

Example request

```
{
  "create_destination": {
    "storage_service": {
      "name": "string"
    },
    "tiering": {
      "policy": "string"
    }
  },
  "destination": {
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ],
    "ipspace": "Default",
    "path": "svm1:volume1",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "exported_snapshot": "string",
  "lag_time": "PT8H35M42S",
  "policy": {
    "name": "Asynchronous",
    "type": "string",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "restore_to_snapshot": "string",
  "source": {
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ]
  }
}
```

```

    }
  ],
  "ipspace": "Default",
  "path": "svm1:volume1",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
},
"state": "snapmirrored",
"transfer": {
  "state": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
},
"unhealthy_reason": [
  {
    "code": "6621444",
    "message": "Failed to complete update operation on one or more
item relationships.",
    "parameters": []
  },
  {
    "code": "6621445",
    "message": "Group Update failed",
    "parameters": []
  }
],
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
6620374	Internal error. Failed to get SVM information.
6620478	Internal error. Failed to check SnapMirror capability.
13303819	Could not retrieve SnapMirror policy information.
13303821	Invalid SnapMirror policy UUID.
13303841	This operation is not supported for SnapMirror relationships between these endpoints.
13303852	destination.path provided does not contain "\:\\".
13303853	Restore relationships are not supported for SVM-DR endpoints.
13303868	Create of destination endpoint and SnapMirror relationship failed.
13303869	Creating a destination endpoint is not supported for restore relationships.
13303870	A tiering policy cannot be specified if tiering is not being set to supported.
13303871	Storage service properties cannot be specified if the storage service is not being enabled.
13303872	Specified property requires a later effective cluster version.
13303873	Specifying a state when creating a relationship is only supported when creating a destination endpoint.
13303874	Specified state is not supported when creating this relationship.
13303875	Destination aggregates do not have sufficient space for hosting copies of source volumes.

Error Code	Description
13303876	Destination cluster does not have composite aggregates.
13303877	Source or destination cluster must be specified.
13303878	The specified fields do not match.
13303879	Source cluster name or UUID is needed to provision a destination SVM on the local cluster.
13303880	Source cluster must be remote for provisioning a destination SVM on the local cluster.
13303881	Network validation failed.
13303882	SVM validation failed.
13303883	Encryption is not enabled on the destination cluster.
13303887	Synchronous SnapMirror relationships between FlexGroup volumes are not supported.
13303888	Synchronous SnapMirror relationships require an effective cluster version of 9.5 or later on both the source and destination clusters.
13303889	Asynchronous SnapMirror relationships between FlexGroup volumes require an effective cluster version of 9.5 or later on both the source and destination clusters.
13303890	Asynchronous SnapMirror relationships between FlexVol volumes require an effective cluster version of 9.3, 9.5, or later on both the source and destination clusters.
13303891	Creating a destination endpoint with storage service requires an effective cluster version of 9.7 or later.
13303892	Fetching remote information from the destination cluster failed.
13303893	Updating job description failed.
13303894	Destination volume name is invalid. It must contain the source volume name and have a suffix when creating a destination endpoint on a cluster with an effective cluster version of 9.6 or earlier.
13303895	Operation on the remote destination cluster is not supported.
13303916	FlexGroup volumes are not supported on SnapLock aggregates.
13303918	No suitable destination aggregate type is available.
13303919	Only FabricPool enabled aggregates are available on the destination.

Error Code	Description
13303920	Only SnapLock aggregates are available on the destination. FlexGroup volumes are not supported on SnapLock aggregates.
13303921	Unable to retrieve the SnapMirror capabilities of the destination cluster.
13303922	Specified source SVM is not a data SVM.
13303923	Specified destination SVM is not a data SVM.
13303924	Source SVM has an invalid Snapshot copy policy.
13303925	SnapMirror validation has failed.
13303930	The specified tiering policy is not supported for destination volumes of Synchronous relationships.
13303938	Fetching information from the local cluster failed.
13303939	Could not create an SVM peer relationship.
13303944	An SVM-DR relationship is not supported because the source SVM has CIFS configured and the associated SnapMirror policy has either the identity_preservation property not set or set to exclude_network_and_protocol_config.
13303966	Consistency Group relationships require a policy of type "sync" with a sync_type of "automated_failover".
13303967	Consistency Group volume is not a FlexVol volume.
13303968	Unsupported volume type for the Consistency Group.
13303969	SnapMirror relationships between SVM endpoints and object store endpoints are not supported.
13303970	Unsupported policy type for the Consistency Group.
13303971	SnapMirror relationships between Consistency Group endpoints and object store endpoints are not supported.
13303976	Source or destination SVM is already part of an SVM-DR relation.
13303977	Destination Consistency Group volume UUIDs are not expected while provisioning the destination volumes.
13303978	Number of Consistency Group volume names and UUIDs does not match.
13303979	Number of Consistency Group volumes exceeds the allowed limit.
13303980	Number of source and destination Consistency Group volumes do not match.
13303981	ISCSI or FCP protocol is not configured.

Error Code	Description
13303982	SAN data interface is not configured on the SVM.
13304021	No suitable storage can be found meeting the specified requirements. No FabricPool enabled aggregates are available on the destination.
13304022	No suitable storage can be found meeting the specified requirements. No non-root, non-taken-over, non-SnapLock, non-composite aggregates are available on the destination.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

status

Name	Type	Description
code	string	Status code
message	string	SnapMirror Consistency Group failover status.

snapmirror_consistency_group_failover

SnapMirror Consistency Group failover information. The SnapMirror Consistency Group failover can be a planned or an unplanned operation. Only active SnapMirror Consistency Group failover operation progress can be monitored using this object. In case of an error during the failover operation, the property "consistency_group_failover.error" holds the reason for the error. ONTAP automatically retries any failed SnapMirror Consistency Group failover operation.

Name	Type	Description
status	status	

storage_service

Name	Type	Description
enabled	boolean	This property indicates whether to create the destination endpoint using storage service.
enforce_performance	boolean	Optional property to enforce storage service performance on the destination endpoint. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints.
name	string	<p>Optional property to specify the storage service name for the destination endpoint. This property is considered when the property "create_destination.storage_service.enabled" is set to "true". When the property "create_destination.storage_service.enabled" is set to "true" and the "create_destination.storage_service.name" for the endpoint is not specified, then ONTAP selects the highest storage service available on the cluster to provision the destination endpoint. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints.</p> <ul style="list-style-type: none"> enum: ["extreme", "performance", "value"] Introduced in: 9.6

tiering

Name	Type	Description
policy	string	<p>Optional property to specify the destination endpoint's tiering policy when "create_destination.tiering.supported" is set to "true". This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. This property determines whether the user data blocks of the destination endpoint in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of the destination endpoint volume blocks increases if they are accessed frequently and decreases when they are not.</p> <p>all – This policy allows tiering of both destination endpoint Snapshot copies and the user transferred data blocks to the cloud store as soon as possible by ignoring the temperature on the volume blocks. This tiering policy is not applicable for Consistency Group destination endpoints or for synchronous relationships.</p> <p>auto – This policy allows tiering of both destination endpoint Snapshot copies and the active file system user data to the cloud store</p> <p>none – Destination endpoint volume blocks will not be tiered to the cloud store.</p> <p>snapshot_only – This policy allows tiering of only the destination endpoint volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot_only" for a FlexVol volume and "none" for a FlexGroup volume.</p>

Name	Type	Description
supported	boolean	Optional property to enable provisioning of the destination endpoint volumes on FabricPool aggregates. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. Only FabricPool aggregates are used if this property is set to "true" and only non FabricPool aggregates are used if this property is set to "false". Tiering support for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only do the destination endpoint volumes need to support tiering by using FabricPools, the "create_destination.tiering.policy" must not be "none". A destination endpoint that uses FabricPools but has a tiering "policy" of "none" supports tiering but will not tier any data.

snapmirror_destination_creation

Use this object to provision the destination endpoint when establishing a SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM, or Consistency Group. Given a source endpoint, the destination endpoint is provisioned in the SVM specified in the "destination.path" property. While protecting an SVM, the SVM destination endpoint can only be provisioned on the local cluster. To provision the SVM destination endpoint use the optional "source.cluster.name" property to specify the remote cluster name or use the optional "source.cluster.uuid" property to specify the remote cluster UUID. When "create_destination.enabled" option is specified while making a POST for a SnapMirror relationship, the relationship can be automatically initialized by setting the "state" either to "snapmirrored" when the policy is of type "async" or to "in_sync" when the policy is of type "sync". The "destination.path" property must specify the destination endpoint path. For example, for FlexVol volume and FlexGroup volume, the "destination.path" can be specified as <destination-svm-name:dp-volume-name>, for SVM data protection, the "destination.path" must be specified as <destination-svm-name:>, and for Consistency Group, the "destination.path" must be specified as <destination-svm-name:> along with the "destination.consistency_group_volumes" property to indicate the list of destination volumes of type "DP" in the Consistency Group. For a FlexVol volume, a FlexGroup volume, or a Consistency Group destination endpoint, the properties in this object can be specified either from the source or the destination cluster. For an SVM destination endpoint, the properties in this object can be specified from the destination cluster only. This object is not supported for non ONTAP endpoints.</destination-svm-name:></destination-svm-name:></destination-svm-name:dp-volume-name>

Name	Type	Description
enabled	boolean	Optional property to create the destination endpoint when establishing a SnapMirror relationship. It is assumed to be "false" if no other property is set and assumed to be "true" if any other property is set.
storage_service	storage_service	
tiering	tiering	

cluster

Name	Type	Description
name	string	
uuid	string	

consistency_group_volumes

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are

peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
cluster	cluster	
consistency_group_volumes	array[consistency_group_volumes]	Mandatory property for a Consistency Group endpoint. Specifies the list of FlexVol volumes for a Consistency Group.
ipspace	string	Optional property to specify the IPspace of the SVM.
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: ONTAP Consistency Group - svm1:/cg/cg_name <ul style="list-style-type: none">• example: svm1:volume1• Introduced in: 9.6
svm	svm	

policy

Basic policy information of the relationship.

Name	Type	Description
name	string	
type	string	
uuid	string	

transfer

Basic information on the current transfer.

Name	Type	Description
bytes_transferred	integer	Bytes transferred.
state	string	
uuid	string	

snapmirror_error

SnapMirror error

Name	Type	Description
code	integer	Error code
message	string	Error message
parameters	array[string]	Parameters for the error message

snapmirror_relationship

SnapMirror relationship information. The SnapMirror relationship can be either "async" or "sync" based on the type of SnapMirror policy associated with the relationship. The source and destination endpoints of a SnapMirror relationship must be of the same type, for example, if the source endpoint is a FlexVol volume then the destination endpoint must be a FlexVol volume.

The SnapMirror policy type "async" can be used when the SnapMirror relationship has FlexVol volume or FlexGroup volume or SVM as the endpoint. The SnapMirror policy type "sync" can be used when the SnapMirror relationship has FlexVol volume as the endpoint. The SnapMirror policy type "sync" with "sync_type" as "automated_failover" can be used when the SnapMirror relationship has Consistency Group as the endpoint.



Name	Type	Description
create_destination	snapmirror_destination_creation	<p>Use this object to provision the destination endpoint when establishing a SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM, or Consistency Group. Given a source endpoint, the destination endpoint is provisioned in the SVM specified in the "destination.path" property. While protecting an SVM, the SVM destination endpoint can only be provisioned on the local cluster. To provision the SVM destination endpoint use the optional "source.cluster.name" property to specify the remote cluster name or use the optional "source.cluster.uuid" property to specify the remote cluster UUID. When "create_destination.enabled" option is specified while making a POST for a SnapMirror relationship, the relationship can be automatically initialized by setting the "state" either to "snapmirrored" when the policy is of type "async" or to "in_sync" when the policy is of type "sync". The "destination.path" property must specify the destination endpoint path. For example, for FlexVol volume and FlexGroup volume, the "destination.path" can be specified as <destination-SVM-name:dp-volume-name>, for SVM data protection, the "destination.path" must be specified as &lt;destination-SVM-name:&gt;, and for Consistency Group, the "destination.path" must be specified as &lt;destination-SVM-name:/cg/consistency-group-name&gt; along with the "destination.consistency_group_volumes" property to indicate the list of destination volumes of type "DP" in the Consistency Group. For a FlexVol volume, a FlexGroup volume, or a Consistency Group destination endpoint, the properties in this object can be specified either from the source or the destination</p>

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
exported_snapshot	string	Snapshot copy exported to clients on destination.
healthy	boolean	Is the relationship healthy?
lag_time	string	Time since the exported Snapshot copy was created.
policy	policy	Basic policy information of the relationship.
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This property is applicable only for relationships with FlexVol volume or FlexGroup volume endpoints and when the PATCH state is being changed to "snapmirrored".

Name	Type	Description
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This property is applicable only for relationships with FlexVol volume endpoints and when the PATCH state is being changed to "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup volume relationship. This restores all destination FlexGroup constituent volumes to the latest Snapshot copy, and any writes to the read-write constituents are lost. This property is applicable only for SnapMirror relationships with FlexGroup volume endpoints and when the PATCH state is being changed to "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship. SnapMirror relationships with the policy type "async" can be restored. SnapMirror relationships with the policy type "sync" cannot be restored.
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination during the break operation. This property is applicable only for SnapMirror relationships with FlexVol volume endpoints and when the PATCH state is being changed to "broken_off".

Name	Type	Description
source	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
state	string	<p>State of the relationship.</p> <p>To initialize the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>To break the relationship, PATCH the state to "broken_off" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "broken_off".</p> <p>To resync the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" can be in "broken_off" state due to a failed attempt of SnapMirror failover.</p> <p>To pause the relationship, suspending further transfers, PATCH the state to "paused" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "paused".</p> <p>To resume transfers for a paused relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships with a policy of type "sync". A PATCH call on the state change only triggers the transition to the specified state. You must poll on the "state", "healthy" and "unhealthy_reason" properties using a GET request to determine if the transition is successful. To automatically initialize the relationship when specifying</p>

Name	Type	Description
transfer	transfer	Basic information on the current transfer.
unhealthy_reason	array[snapmirror_error]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

Delete a SnapMirror relationship

DELETE /snapmirror/relationships/{uuid}

Introduced In: 9.6

Deletes a SnapMirror relationship.

Important notes

- The "destination_only", "source_only", and "source_info_only" flags are mutually exclusive. If no flag is specified, the relationship is deleted from both the source and destination and all common Snapshot copies between the source and destination are also deleted.
- For a restore relationship, the call must be executed on the cluster containing the destination endpoint without specifying the destination_only, source_only, or source_info_only parameters.
- Additionally, ensure that there are no ongoing transfers on a restore relationship before calling this API.
- The "failover", "force-failover" and "failback" query parameters are only applicable for SVM-DR SnapMirror relationships.

Related ONTAP commands

- `snapmirror delete`
- `snapmirror release`

Examples

The following examples show how to delete the relationship from both the source and destination, the destination only, and the source only.

Deleting the relationship from both the source and destination. This API must be run on the cluster containing

the destination endpoint.

```
DELETE "/api/snapmirror/relationships/4512b2d2-fd60-11e8-8929-005056bbfe52"
```

Deleting the relationship on the destination only. This API must be run on the cluster containing the destination endpoint.

```
DELETE "/api/snapmirror/relationships/fd1e0697-02ba-11e9-acc7-005056a7697f/?destination_only=true"
```

Deleting the relationship on the source only. This API must be run on the cluster containing the source endpoint.

```
DELETE "/api/snapmirror/relationships/93e828ba-02bc-11e9-acc7-005056a7697f/?source_only=true"
```

Deleting the source information only. This API must be run on the cluster containing the source endpoint. This does not delete the common Snapshot copies between the source and destination.

```
DELETE "/api/snapmirror/relationships/caf545a2-fc60-11e8-aa13-005056a707ff/?source_info_only=true"
```

Learn more

- [DOC /snapmirror/relationships](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Relationship UUID
destination_only	boolean	query	False	Deletes a relationship on the destination only. This parameter is applicable only when the call is executed on the cluster that contains the destination endpoint.

Name	Type	In	Required	Description
source_only	boolean	query	False	Deletes a relationship on the source only. This parameter is applicable only when the call is executed on the cluster that contains the source endpoint.
source_info_only	boolean	query	False	Deletes relationship information on the source only. This parameter is applicable only when the call is executed on the cluster that contains the source endpoint.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303825	Could not retrieve information for the SnapMirror policy type
13303814	Could not retrieve the source or destination SVM UUID
13303815	Could not retrieve information for the peer cluster
13303822	SnapMirror release has failed
13303813	SnapMirror release was successful but delete has failed
13303854	Cleanup of restore relationship failed
13303855	DELETE call on a restore relationship does not support the given flags
13303865	Deleting the specified SnapMirror policy is not supported.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a SnapMirror relationship

GET /snapmirror/relationships/{uuid}

Introduced In: 9.6

Retrieves a SnapMirror relationship.

Related ONTAP commands

- `snapmirror show`
- `snapmirror list-destinations`

Example

```
GET "/api/snapmirror/relationships/caf545a2-fc60-11e8-aa13-005056a707ff/"
```

Learn more

- [DOC /snapmirror/relationships](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Relationship UUID
list_destinations_only	boolean	query	False	Set to true to show relationships from the source only.
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	

Name	Type	Description
consistency_group_failover	snapmirror_consistency_group_failover	<p>SnapMirror Consistency Group failover information. The SnapMirror Consistency Group failover can be a planned or an unplanned operation. Only active SnapMirror Consistency Group failover operation progress can be monitored using this object. In case of an error during the failover operation, the property "consistency_group_failover.error" holds the reason for the error. ONTAP automatically retries any failed SnapMirror Consistency Group failover operation.</p> <ul style="list-style-type: none"> • Introduced in: 9.8 • readOnly: 1
destination	snapmirror_endpoint	<p>Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.</p>
exported_snapshot	string	Snapshot copy exported to clients on destination.
healthy	boolean	Is the relationship healthy?

Name	Type	Description
lag_time	string	Time since the exported Snapshot copy was created.
policy	policy	Basic policy information of the relationship.
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This property is applicable only for relationships with FlexVol volume or FlexGroup volume endpoints and when the PATCH state is being changed to "snapmirrored".
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This property is applicable only for relationships with FlexVol volume endpoints and when the PATCH state is being changed to "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup volume relationship. This restores all destination FlexGroup constituent volumes to the latest Snapshot copy, and any writes to the read-write constituents are lost. This property is applicable only for SnapMirror relationships with FlexGroup volume endpoints and when the PATCH state is being changed to "broken_off".
restore	boolean	Set to true to create a relationship for restore. To trigger restore-transfer, use transfers POST on the restore relationship. SnapMirror relationships with the policy type "async" can be restored. SnapMirror relationships with the policy type "sync" cannot be restored.

Name	Type	Description
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination during the break operation. This property is applicable only for SnapMirror relationships with FlexVol volume endpoints and when the PATCH state is being changed to "broken_off".
source	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
state	string	<p>State of the relationship.</p> <p>To initialize the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>To break the relationship, PATCH the state to "broken_off" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "broken_off".</p> <p>To resync the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" can be in "broken_off" state due to a failed attempt of SnapMirror failover.</p> <p>To pause the relationship, suspending further transfers, PATCH the state to "paused" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "paused".</p> <p>To resume transfers for a paused relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships with a policy of type "sync". A PATCH call on the state change only triggers the transition to the specified state. You must poll on the "state", "healthy" and "unhealthy_reason" properties using a GET request to determine if the transition is successful. To automatically initialize the relationship when specifying "create_destination" property, set the state to "snapmirrored" for</p>

Name	Type	Description
transfer	transfer	Basic information on the current transfer.
unhealthy_reason	array[snapmirror_error]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "consistency_group_failover": {
    "error": {
      "arguments": [
        {
          "code": "string",
          "message": "string"
        }
      ],
      "code": "4",
      "message": "entry doesn't exist",
      "target": "uuid"
    },
    "status": {
      "code": "string",
      "message": "string"
    }
  },
  "destination": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ]
  },
}
```

```

    "path": "svm1:volume1",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "exported_snapshot": "string",
  "lag_time": "PT8H35M42S",
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "Asynchronous",
    "type": "string",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "restore_to_snapshot": "string",
  "source": {
    "cluster": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ]
  },
  "path": "svm1:volume1",

```

```

    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "state": "snapmirrored",
    "transfer": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "state": "string",
      "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
    },
    "unhealthy_reason": [
      {
        "code": "6621444",
        "message": "Failed to complete update operation on one or more
item relationships.",
        "parameters": []
      },
      {
        "code": "6621445",
        "message": "Group Update failed",
        "parameters": []
      }
    ],
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  }
}

```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303825	Could not retrieve information for the SnapMirror policy type
13303817	Unknown value for the Snapmirror State

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

status

Name	Type	Description
code	string	Status code
message	string	SnapMirror Consistency Group failover status.

snapmirror_consistency_group_failover

SnapMirror Consistency Group failover information. The SnapMirror Consistency Group failover can be a planned or an unplanned operation. Only active SnapMirror Consistency Group failover operation progress can be monitored using this object. In case of an error during the failover operation, the property "consistency_group_failover.error" holds the reason for the error. ONTAP automatically retries any failed

SnapMirror Consistency Group failover operation.

Name	Type	Description
error	error	
status	status	

storage_service

Name	Type	Description
enabled	boolean	This property indicates whether to create the destination endpoint using storage service.
enforce_performance	boolean	Optional property to enforce storage service performance on the destination endpoint. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints.
name	string	<p>Optional property to specify the storage service name for the destination endpoint. This property is considered when the property "create_destination.storage_service.enabled" is set to "true". When the property "create_destination.storage_service.enabled" is set to "true" and the "create_destination.storage_service.name" for the endpoint is not specified, then ONTAP selects the highest storage service available on the cluster to provision the destination endpoint. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints.</p> <ul style="list-style-type: none">enum: ["extreme", "performance", "value"]Introduced in: 9.6

tiering

Name	Type	Description
policy	string	<p>Optional property to specify the destination endpoint's tiering policy when "create_destination.tiering.supported" is set to "true". This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. This property determines whether the user data blocks of the destination endpoint in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of the destination endpoint volume blocks increases if they are accessed frequently and decreases when they are not.</p> <p>all – This policy allows tiering of both destination endpoint Snapshot copies and the user transferred data blocks to the cloud store as soon as possible by ignoring the temperature on the volume blocks. This tiering policy is not applicable for Consistency Group destination endpoints or for synchronous relationships.</p> <p>auto – This policy allows tiering of both destination endpoint Snapshot copies and the active file system user data to the cloud store</p> <p>none – Destination endpoint volume blocks will not be tiered to the cloud store.</p> <p>snapshot_only – This policy allows tiering of only the destination endpoint volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot_only" for a FlexVol volume and "none" for a FlexGroup volume.</p>

Name	Type	Description
supported	boolean	Optional property to enable provisioning of the destination endpoint volumes on FabricPool aggregates. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. Only FabricPool aggregates are used if this property is set to "true" and only non FabricPool aggregates are used if this property is set to "false". Tiering support for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only do the destination endpoint volumes need to support tiering by using FabricPools, the "create_destination.tiering.policy" must not be "none". A destination endpoint that uses FabricPools but has a tiering "policy" of "none" supports tiering but will not tier any data.

snapmirror_destination_creation

Use this object to provision the destination endpoint when establishing a SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM, or Consistency Group. Given a source endpoint, the destination endpoint is provisioned in the SVM specified in the "destination.path" property. While protecting an SVM, the SVM destination endpoint can only be provisioned on the local cluster. To provision the SVM destination endpoint use the optional "source.cluster.name" property to specify the remote cluster name or use the optional "source.cluster.uuid" property to specify the remote cluster UUID. When "create_destination.enabled" option is specified while making a POST for a SnapMirror relationship, the relationship can be automatically initialized by setting the "state" either to "snapmirrored" when the policy is of type "async" or to "in_sync" when the policy is of type "sync". The "destination.path" property must specify the destination endpoint path. For example, for FlexVol volume and FlexGroup volume, the "destination.path" can be specified as <destination-svm-name:dp-volume-name>, for SVM data protection, the "destination.path" must be specified as <destination-svm-name:>, and for Consistency Group, the "destination.path" must be specified as <destination-svm-name:> along with the "destination.consistency_group_volumes" property to indicate the list of destination volumes of type "DP" in the Consistency Group. For a FlexVol volume, a FlexGroup volume, or a Consistency Group destination endpoint, the properties in this object can be specified either from the source or the destination cluster. For an SVM destination endpoint, the properties in this object can be specified from the destination cluster only. This object is not supported for non ONTAP endpoints.</destination-svm-name:></destination-svm-name:></destination-svm-name:dp-volume-name>

Name	Type	Description
enabled	boolean	Optional property to create the destination endpoint when establishing a SnapMirror relationship. It is assumed to be "false" if no other property is set and assumed to be "true" if any other property is set.
storage_service	storage_service	
tiering	tiering	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

consistency_group_volumes

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the

endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
cluster	cluster	
consistency_group_volumes	array[consistency_group_volumes]	Mandatory property for a Consistency Group endpoint. Specifies the list of FlexVol volumes for a Consistency Group.
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: ONTAP Consistency Group - svm1:/cg/cg_name <ul style="list-style-type: none"> • example: svm1:volume1 • Introduced in: 9.6
svm	svm	

policy

Basic policy information of the relationship.

Name	Type	Description
_links	_links	
name	string	
type	string	
uuid	string	

transfer

Basic information on the current transfer.

Name	Type	Description
_links	_links	
bytes_transferred	integer	Bytes transferred.
state	string	

Name	Type	Description
uuid	string	

snapmirror_error

SnapMirror error

Name	Type	Description
code	integer	Error code
message	string	Error message
parameters	array[string]	Parameters for the error message

Update a SnapMirror relationship

PATCH /snapmirror/relationships/{uuid}

Introduced In: 9.6

Updates a SnapMirror relationship. This API is used to initiate SnapMirror operations such as "initialize", "resync", "break", "quiesce", and "resume" by specifying the appropriate value for the "state" field. It is also used to modify the SnapMirror policy associated with the specified relationship. Additionally, a SnapMirror relationship can be failed over to the destination endpoint or a failed over SnapMirror relationship can be failed back to the original state or a SnapMirror relationship direction can be reversed using this API.

To initialize the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or "in_sync" for relationships with a policy of type "sync".

To break the relationship or to failover to the destination endpoint and start serving data from the destination endpoint, PATCH the state to "broken_off" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and sync_type as "automated_failover" cannot be "broken_off".

To resync the broken relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or "in_sync" for relationships with a policy of type "sync".

To failback the failed over relationship and start serving data from the source endpoint, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or "in_sync" for relationships with a policy of type "sync" and set the query flag "failback" as "true". SnapMirror relationships with the policy type as "sync" and sync_type as "automated_failover" can be in "broken_off" state due to a failed attempt of automated SnapMirror failover operation.

To pause the relationship, suspending further transfers, PATCH the state to "paused" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and sync_type as "automated_failover" cannot be "paused".

To resume transfers for a paused relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or "in_sync" for relationships with a policy of type "sync".

To reverse the direction of the relationship, PATCH the "source.path" with the destination endpoint and the

"destination.path" with the source endpoint and the relationship state to "snapmirrored" for relationships with a policy of type "async" or "in_sync" for relationships with a policy of type "sync".

The values "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships with a policy of type "sync".

Related ONTAP commands

- `snapmirror modify`
- `snapmirror initialize`
- `snapmirror resync`
- `snapmirror break`
- `snapmirror quiesce`
- `snapmirror resume`

Examples

The following examples show how to perform the SnapMirror "resync", "initialize", "resume", "quiesce", and "break" operations. In addition, a relationship can be failed over to the destination endpoint and start serving data from the destination endpoint. A failed over relationship can be failed back to the source endpoint and serve data from the source endpoint. Also a relationship can be reversed by making the source endpoint as the new destination endpoint and the destination endpoint as the new source endpoint.

To update an associated SnapMirror policy.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"policy": { "name" : "MirrorAndVaultDiscardNetwork"}}'
```

To perform SnapMirror "resync" for an asynchronous SnapMirror relationship.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"state":"snapmirrored"}'
```

To perform SnapMirror "initialize" for an asynchronous SnapMirror relationship.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"state":"snapmirrored"}'
```

To perform SnapMirror "resume" for an asynchronous SnapMirror relationship.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"state":"snapmirrored"}'
```

To perform SnapMirror "quiesce" for an asynchronous SnapMirror relationship.


```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff" '{"state":"paused"}'
```

To perform SnapMirror "break" for an asynchronous SnapMirror relationship. This operation does a failover to the destination endpoint. After a the failover, data can then be served from the destination endpoint.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff" '{"state":"broken_off"}'
```

To forcefully failover to the destination endpoint and start serving data from the destination endpoint.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/?force=true" '{"state":"broken_off"}'
```

To failback to the source endpoint and start serving data from the source endpoint for an asynchronous relationship.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/?failback=true" '{"state":"snapmirrored"}'
```

To failback to the source endpoint and start serving data from the source endpoint for a synchronous relationship.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/?failback=true" '{"state":"in_sync"}'
```

To reverse the direction of an asynchronous relationship, that is, make the source endpoint as the new destination endpoint and make the destination endpoint as the new source endpoint.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"source": {"path": "dst_svm:dst_vol"}, "destination": {"path": "src_svm:src_vol"}, "state": "snapmirrored"}'
```

To reverse the direction of a synchronous relationship, that is, make the source endpoint as the new destination endpoint and make the destination endpoint as the new source endpoint.

```
PATCH "/api/snapmirror/relationships/98bb2608-fc60-11e8-aa13-005056a707ff/" '{"source": {"path": "dst_svm:dst_vol"}, "destination": {"path": "src_svm:src_vol"}, "state": "in_sync"}'
```

Learn more

- [DOC /snapmirror/relationships](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Relationship UUID
failover	boolean	query	False	<p>If this parameter is set, validation and failover will occur to the SVM-DR SnapMirror relationship destination endpoint. Any other fields specified with this parameter will be ignored. This parameter is supported only for SVM-DR SnapMirror relationships.</p> <ul style="list-style-type: none">• Introduced in: 9.7• Default value:
force-failover	boolean	query	False	<p>If this parameter is set, failover will occur to the SVM-DR SnapMirror relationship destination endpoint, overriding the validation errors. Any other fields specified with this parameter will be ignored. This parameter is supported only for SVM-DR SnapMirror relationships.</p> <ul style="list-style-type: none">• Introduced in: 9.7• Default value:

Name	Type	In	Required	Description
force	boolean	query	False	<p>If this parameter is set while specifying the state as "broken_off", indicates a forced failover overriding the validation errors.</p> <ul style="list-style-type: none"> • Introduced in: 9.8 • Default value:
failback	boolean	query	False	<p>If this parameter is set while specifying the state as "snapmirrored", indicates recovery of the failed over SnapMirror relationship by preserving the data written on the destination endpoint when the SnapMirror relationship was in failed over state. This flag is only applicable to SVM-DR SnapMirror relationships.</p> <ul style="list-style-type: none"> • Introduced in: 9.8 • Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
validate_only	boolean	query	False	<p>Validate the operation and its parameters, without actually performing the operation.</p> <ul style="list-style-type: none"> • Introduced in: 9.7

Request Body

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
exported_snapshot	string	Snapshot copy exported to clients on destination.
healthy	boolean	Is the relationship healthy?
lag_time	string	Time since the exported Snapshot copy was created.
policy	policy	Basic policy information of the relationship.
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This property is applicable only for relationships with FlexVol volume or FlexGroup volume endpoints and when the PATCH state is being changed to "snapmirrored".

Name	Type	Description
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This property is applicable only for relationships with FlexVol volume endpoints and when the PATCH state is being changed to "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup volume relationship. This restores all destination FlexGroup constituent volumes to the latest Snapshot copy, and any writes to the read-write constituents are lost. This property is applicable only for SnapMirror relationships with FlexGroup volume endpoints and when the PATCH state is being changed to "broken_off".
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination during the break operation. This property is applicable only for SnapMirror relationships with FlexVol volume endpoints and when the PATCH state is being changed to "broken_off".

Name	Type	Description
source	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
state	string	<p>State of the relationship.</p> <p>To initialize the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>To break the relationship, PATCH the state to "broken_off" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "broken_off".</p> <p>To resync the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" can be in "broken_off" state due to a failed attempt of SnapMirror failover.</p> <p>To pause the relationship, suspending further transfers, PATCH the state to "paused" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "paused".</p> <p>To resume transfers for a paused relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships with a policy of type "sync". A PATCH call on the state change only triggers the transition to the specified state. You must poll on the "state", "healthy" and "unhealthy_reason" properties using a GET request to determine if the transition is successful. To automatically initialize the relationship when specifying "create_destination" property, set the state to "snapmirrored" for</p>

Name	Type	Description
transfer	transfer	Basic information on the current transfer.
unhealthy_reason	array[snapmirror_error]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

Example request

```
{
  "destination": {
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ],
    "path": "svm1:volume1",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
  "exported_snapshot": "string",
  "lag_time": "PT8H35M42S",
  "policy": {
    "name": "Asynchronous",
    "type": "string",
    "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
  },
  "restore": null,
  "restore_to_snapshot": "string",
  "source": {
    "cluster": {
      "name": "cluster1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "consistency_group_volumes": [
      {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    ],
    "path": "svm1:volume1",
    "svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    }
  },
}
```

```

"state": "snapmirrored",
"transfer": {
  "state": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
},
"unhealthy_reason": [
  {
    "code": "6621444",
    "message": "Failed to complete update operation on one or more
item relationships.",
    "parameters": []
  },
  {
    "code": "6621445",
    "message": "Group Update failed",
    "parameters": []
  }
],
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

{
  "job": {
    "uuid": "string"
  }
}

```

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303825	Could not retrieve information for the SnapMirror policy type
13303817	Unknown value for the SnapMirror state
13303829	Invalid state
13303830	Transient state
13303831	Invalid state for async SnapMirror relationship
13303834	Given input valid only for FlexGroup SnapMirror relationship
13303835	Given flag is valid only when PATCH state is broken_off
13303836	Given flag is valid only when PATCH state is snapmirrored or in_sync
13303818	Invalid state transition requested
13303828	Given state change is not possible for SVM SnapMirror relationship
13303833	Requested state change is not possible
13303832	SnapMirror relationship is already initialized
13303824	Quiescing the SnapMirror relationship has failed
13303826	Required environment variables are not set
13303827	Internal Error
13303823	Quiesce operation timed out
13303821	Invalid SnapMirror policy name/UUID
13303819	Could not retrieve SnapMirror policy information
13303851	Cannot modify attributes of SnapMirror restore relationship
13303816	Could not retrieve state or status values
13303837	Given flags are valid only if SnapMirror state change is requested
6619546	Destination must be a dp volume
13303808	Transition to broken_off state failed
13303809	Transition to paused state failed
13303810	Transition to snapmirrored state failed
13303811	Transition from paused state failed
13303820	SnapMirror policy was successfully updated, state transition failed
13303856	SVM is not configured with any data protocol

Error code	Description
13303857	SVM is not configured with any network interface
13303858	Internal error. Failed to check LIF and protocols details for SVM
13303859	Internal error. SVM Failover operation failed. SVM operational state is unavailable.
13303865	Modifying the specified SnapMirror policy is not supported.
13303866	Cannot use the specified policy to modify the policy of the relationship.
13303867	Modifying the policy of an async-mirror or a vault relationship is not supported.
13303884	LIF and protocols details are configured incorrectly for SVM.
13303996	The source and destination clusters both have a policy with the same name, but they have different properties.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

status

Name	Type	Description
code	string	Status code
message	string	SnapMirror Consistency Group failover status.

snapmirror_consistency_group_failover

SnapMirror Consistency Group failover information. The SnapMirror Consistency Group failover can be a planned or an unplanned operation. Only active SnapMirror Consistency Group failover operation progress can be monitored using this object. In case of an error during the failover operation, the property "consistency_group_failover.error" holds the reason for the error. ONTAP automatically retries any failed SnapMirror Consistency Group failover operation.

Name	Type	Description
status	status	

storage_service

Name	Type	Description
enabled	boolean	This property indicates whether to create the destination endpoint using storage service.
enforce_performance	boolean	Optional property to enforce storage service performance on the destination endpoint. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints.
name	string	<p>Optional property to specify the storage service name for the destination endpoint. This property is considered when the property "create_destination.storage_service.enabled" is set to "true". When the property "create_destination.storage_service.enabled" is set to "true" and the "create_destination.storage_service.name" for the endpoint is not specified, then ONTAP selects the highest storage service available on the cluster to provision the destination endpoint. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints.</p> <ul style="list-style-type: none"> enum: ["extreme", "performance", "value"] Introduced in: 9.6

tiering

Name	Type	Description
policy	string	<p>Optional property to specify the destination endpoint's tiering policy when "create_destination.tiering.supported" is set to "true". This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. This property determines whether the user data blocks of the destination endpoint in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of the destination endpoint volume blocks increases if they are accessed frequently and decreases when they are not.</p> <p>all – This policy allows tiering of both destination endpoint Snapshot copies and the user transferred data blocks to the cloud store as soon as possible by ignoring the temperature on the volume blocks. This tiering policy is not applicable for Consistency Group destination endpoints or for synchronous relationships.</p> <p>auto – This policy allows tiering of both destination endpoint Snapshot copies and the active file system user data to the cloud store</p> <p>none – Destination endpoint volume blocks will not be tiered to the cloud store.</p> <p>snapshot_only – This policy allows tiering of only the destination endpoint volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot_only" for a FlexVol volume and "none" for a FlexGroup volume.</p>

Name	Type	Description
supported	boolean	Optional property to enable provisioning of the destination endpoint volumes on FabricPool aggregates. This property is applicable to FlexVol volume, FlexGroup volume, and Consistency Group endpoints. Only FabricPool aggregates are used if this property is set to "true" and only non FabricPool aggregates are used if this property is set to "false". Tiering support for a FlexGroup volume can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only do the destination endpoint volumes need to support tiering by using FabricPools, the "create_destination.tiering.policy" must not be "none". A destination endpoint that uses FabricPools but has a tiering "policy" of "none" supports tiering but will not tier any data.

snapmirror_destination_creation

Use this object to provision the destination endpoint when establishing a SnapMirror relationship for a FlexVol volume, FlexGroup volume, SVM, or Consistency Group. Given a source endpoint, the destination endpoint is provisioned in the SVM specified in the "destination.path" property. While protecting an SVM, the SVM destination endpoint can only be provisioned on the local cluster. To provision the SVM destination endpoint use the optional "source.cluster.name" property to specify the remote cluster name or use the optional "source.cluster.uuid" property to specify the remote cluster UUID. When "create_destination.enabled" option is specified while making a POST for a SnapMirror relationship, the relationship can be automatically initialized by setting the "state" either to "snapmirrored" when the policy is of type "async" or to "in_sync" when the policy is of type "sync". The "destination.path" property must specify the destination endpoint path. For example, for FlexVol volume and FlexGroup volume, the "destination.path" can be specified as <destination-svm-name:dp-volume-name>, for SVM data protection, the "destination.path" must be specified as <destination-svm-name:>, and for Consistency Group, the "destination.path" must be specified as <destination-svm-name:> along with the "destination.consistency_group_volumes" property to indicate the list of destination volumes of type "DP" in the Consistency Group. For a FlexVol volume, a FlexGroup volume, or a Consistency Group destination endpoint, the properties in this object can be specified either from the source or the destination cluster. For an SVM destination endpoint, the properties in this object can be specified from the destination cluster only. This object is not supported for non ONTAP endpoints.</destination-svm-name:></destination-svm-name:></destination-svm-name:dp-volume-name>

Name	Type	Description
enabled	boolean	Optional property to create the destination endpoint when establishing a SnapMirror relationship. It is assumed to be "false" if no other property is set and assumed to be "true" if any other property is set.
storage_service	storage_service	
tiering	tiering	

cluster

Name	Type	Description
name	string	
uuid	string	

consistency_group_volumes

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are

peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
cluster	cluster	
consistency_group_volumes	array[consistency_group_volumes]	Mandatory property for a Consistency Group endpoint. Specifies the list of FlexVol volumes for a Consistency Group.
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: ONTAP Consistency Group - svm1:/cg/cg_name <ul style="list-style-type: none">• example: svm1:volume1• Introduced in: 9.6
svm	svm	

policy

Basic policy information of the relationship.

Name	Type	Description
name	string	
type	string	
uuid	string	

transfer

Basic information on the current transfer.

Name	Type	Description
bytes_transferred	integer	Bytes transferred.
state	string	
uuid	string	

snapmirror_error

SnapMirror error

Name	Type	Description
code	integer	Error code
message	string	Error message
parameters	array[string]	Parameters for the error message

snapmirror_relationship

SnapMirror relationship information. The SnapMirror relationship can be either "async" or "sync" based on the type of SnapMirror policy associated with the relationship. The source and destination endpoints of a SnapMirror relationship must be of the same type, for example, if the source endpoint is a FlexVol volume then the destination endpoint must be a FlexVol volume.

The SnapMirror policy type "async" can be used when the SnapMirror relationship has FlexVol volume or FlexGroup volume or SVM as the endpoint. The SnapMirror policy type "sync" can be used when the SnapMirror relationship has FlexVol volume as the endpoint. The SnapMirror policy type "sync" with "sync_type" as "automated_failover" can be used when the SnapMirror relationship has Consistency Group as the endpoint.

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
exported_snapshot	string	Snapshot copy exported to clients on destination.

Name	Type	Description
healthy	boolean	Is the relationship healthy?
lag_time	string	Time since the exported Snapshot copy was created.
policy	policy	Basic policy information of the relationship.
preserve	boolean	Set to true on resync to preserve Snapshot copies on the destination that are newer than the latest common Snapshot copy. This property is applicable only for relationships with FlexVol volume or FlexGroup volume endpoints and when the PATCH state is being changed to "snapmirrored".
quick_resync	boolean	Set to true to reduce resync time by not preserving storage efficiency. This property is applicable only for relationships with FlexVol volume endpoints and when the PATCH state is being changed to "snapmirrored".
recover_after_break	boolean	Set to true to recover from a failed SnapMirror break operation on a FlexGroup volume relationship. This restores all destination FlexGroup constituent volumes to the latest Snapshot copy, and any writes to the read-write constituents are lost. This property is applicable only for SnapMirror relationships with FlexGroup volume endpoints and when the PATCH state is being changed to "broken_off".
restore_to_snapshot	string	Specifies the Snapshot copy to restore to on the destination during the break operation. This property is applicable only for SnapMirror relationships with FlexVol volume endpoints and when the PATCH state is being changed to "broken_off".

Name	Type	Description
source	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
state	string	<p>State of the relationship.</p> <p>To initialize the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>To break the relationship, PATCH the state to "broken_off" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "broken_off".</p> <p>To resync the relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" can be in "broken_off" state due to a failed attempt of SnapMirror failover.</p> <p>To pause the relationship, suspending further transfers, PATCH the state to "paused" for relationships with a policy of type "async" or "sync". SnapMirror relationships with the policy type as "sync" and "sync_type" as "automated_failover" cannot be "paused".</p> <p>To resume transfers for a paused relationship, PATCH the state to "snapmirrored" for relationships with a policy of type "async" or to state "in_sync" for relationships with a policy of type "sync".</p> <p>The entries "in_sync", "out_of_sync", and "synchronizing" are only applicable to relationships with a policy of type "sync". A PATCH call on the state change only triggers the transition to the specified state. You must poll on the "state", "healthy" and "unhealthy_reason" properties using a GET request to determine if the transition is successful. To automatically initialize the relationship when specifying</p>

Name	Type	Description
transfer	transfer	Basic information on the current transfer.
unhealthy_reason	array[snapmirror_error]	Reason the relationship is not healthy. It is a concatenation of up to four levels of error messages.
uuid	string	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

Manage SnapMirror relationship transfers

SnapMirror relationships relationship.uuid transfers endpoint overview

Overview

This API is used to manage transfers on an existing SnapMirror relationship.

You can initiate SnapMirror operations such as "initialize", "update", "restore-transfer", and "abort" using this API on asynchronous SnapMirror relationship. On a synchronous SnapMirror relationship, you can initiate SnapMirror "initialize" operation. This API only manages the active transfers on the specified relationship. For the restore relationships, the POST on transfers API triggers "restore-transfer". Successful completion of "restore" also deletes the restore relationship. If the "restore" fails, DELETE on relationships must be called to delete the restore relationship.

Retrieve ongoing SnapMirror transfers for a relationship

GET /snapmirror/relationships/{relationship.uuid}/transfers

Introduced In: 9.6

Retrieves the list of ongoing SnapMirror transfers for the specified relationship.

Related ONTAP commands

- `snapmirror show`

Example


```
GET "/api/snapmirror/relationships/293baa53-e63d-11e8-bff1-005056a793dd/transfers"
```

Learn more

- [DOC /snapmirror/relationships/{relationship.uuid}/transfers](#)

Parameters

Name	Type	In	Required	Description
relationship.uuid	string	path	True	Relationship UUID
relationship.uuid	string	query	False	Filter by relationship.uuid
relationship.restore	boolean	query	False	Filter by relationship.restore
relationship.destination.cluster.name	string	query	False	Filter by relationship.destination.cluster.name <ul style="list-style-type: none">• Introduced in: 9.7
relationship.destination.cluster.uuid	string	query	False	Filter by relationship.destination.cluster.uuid <ul style="list-style-type: none">• Introduced in: 9.7
relationship.destination.consistency_group_volumes.uuid	string	query	False	Filter by relationship.destination.consistency_group_volumes.uuid <ul style="list-style-type: none">• Introduced in: 9.8
relationship.destination.consistency_group_volumes.name	string	query	False	Filter by relationship.destination.consistency_group_volumes.name <ul style="list-style-type: none">• Introduced in: 9.8

Name	Type	In	Required	Description
relationship.destination.path	string	query	False	Filter by relationship.destination.path
relationship.destination.svm.uuid	string	query	False	Filter by relationship.destination.svm.uuid
relationship.destination.svm.name	string	query	False	Filter by relationship.destination.svm.name
state	string	query	False	Filter by state
snapshot	string	query	False	Filter by snapshot
uuid	string	query	False	Filter by uuid
checkpoint_size	integer	query	False	Filter by checkpoint_size
bytes_transferred	integer	query	False	Filter by bytes_transferred
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[snapmirror_transfer]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "bytes_transferred": 0,
      "checkpoint_size": 0,
      "relationship": {
        "destination": {
          "cluster": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            }
          },
          "name": "cluster1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "consistency_group_volumes": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "volume1",
            "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
          }
        ],
        "path": "svm1:volume1",
        "svm": {
          "_links": {
            "self": {
```

```

        "href": "/api/resourcelink"
      }
    },
    "name": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
},
"uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
},
"snapshot": "string",
"state": "string",
"storage_efficiency_enabled": null,
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```


See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

files

Specifies a file or LUN consisting of a source_path and an optional destination_path. If not specified, the destination_path is the same as the source_path.

Name	Type	Description
destination_path	string	
source_path	string	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

consistency_group_volumes

Name	Type	Description
_links	_links	
name	string	The name of the volume.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
cluster	cluster	
consistency_group_volumes	array[consistency_group_volumes]	Mandatory property for a Consistency Group endpoint. Specifies the list of FlexVol volumes for a Consistency Group.
path	string	<p>ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: ONTAP Consistency Group - svm1:/cg/cg_name</p> <ul style="list-style-type: none"> example: svm1:volume1 Introduced in: 9.6
svm	svm	

relationship

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
restore	boolean	Is the relationship for restore?
uuid	string	

snapmirror_transfer

SnapMirror transfer information

Name	Type	Description
_links	_links	
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
relationship	relationship	
snapshot	string	Name of Snapshot copy being transferred.

Name	Type	Description
state	string	Status of the transfer. Set PATCH state to "aborted" to abort the transfer. Set PATCH state to "hard_aborted" to abort the transfer and discard the restart checkpoint.
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Start a SnapMirror transfer operation

POST /snapmirror/relationships/{relationship.uuid}/transfers

Introduced In: 9.6

Starts a SnapMirror transfer operation. This API initiates a restore operation if the SnapMirror relationship is of type "restore". Otherwise, it initiates a SnapMirror "initialize" operation or "update" operation based on the current SnapMirror state.

Default property values

- `storage_efficiency_enabled` - *true*

Related ONTAP commands

- `snapmirror update`

- snapmirror initialize
- snapmirror restore

Examples

The following examples show how to perform SnapMirror "initialize", "update", and "restore" operations.

Perform SnapMirror initialize or update

```
POST "/api/snapmirror/relationships/e4e7e130-0279-11e9-b566-0050568e9909/transfers" '{}'
```

Perform SnapMirror restore transfer

```
POST "/api/snapmirror/relationships/c8c62a90-0fef-11e9-b09e-0050568e7067/transfers" '{"source-snapshot": "src", "files": {"source_path": ["/a1.txt.0"], "destination_path": ["/a1-renamed.txt.0"]}]}'
```

Learn more

- [DOC /snapmirror/relationships/{relationship.uuid}/transfers](#)

Parameters

Name	Type	In	Required	Description
relationship.uuid	string	path	True	Relationship UUID
return_records	boolean	query	False	<div>The default is false. If set to true, the records are returned.</div> <div>• Default value:</div>

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
files	array[files]	This is supported for transfer of restore relationship only. This specifies the list of files or LUNs to be restored. Can contain up to eight files or LUNs.

Name	Type	Description
relationship	relationship	
snapshot	string	Name of Snapshot copy being transferred.
source_snapshot	string	Specifies the Snapshot copy on the source to be transferred to the destination.
storage_efficiency_enabled	boolean	This is supported for transfer of restore relationship only. Set this property to "false" to turn off storage efficiency for data transferred over the wire and written to the destination.
uuid	string	

Example request

```
{
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "files": [
    {
      "destination_path": "/dirb/file2",
      "source_path": "/dira/file1"
    }
  ],
  "relationship": {
    "destination": {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "consistency_group_volumes": [
        {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ],
      "ipspace": "Default",
      "path": "svm1:volume1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    },
    "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
  },
  "snapshot": "string",
  "source_snapshot": "string",
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303845	Restore operation failed
13303812	Initialize operation failed
13303844	Update operation failed
13303846	Empty source path file list
13303847	Invalid arguments
6620237	SnapMirror relationship database write failed
6620238	SnapMirror relationship database read failed

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

files

Specifies a file or LUN consisting of a source_path and an optional destination_path. If not specified, the destination_path is the same as the source_path.

Name	Type	Description
destination_path	string	
source_path	string	

cluster

Name	Type	Description
name	string	
uuid	string	

consistency_group_volumes

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
cluster	cluster	
consistency_group_volumes	array[consistency_group_volumes]	Mandatory property for a Consistency Group endpoint. Specifies the list of FlexVol volumes for a Consistency Group.
ipspace	string	Optional property to specify the IPspace of the SVM.
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: ONTAP Consistency Group - svm1:/cg/cg_name <ul style="list-style-type: none"> • example: svm1:volume1 • Introduced in: 9.6
svm	svm	

relationship

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
restore	boolean	Is the relationship for restore?
uuid	string	

snapmirror_transfer

SnapMirror transfer information

Name	Type	Description
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
files	array[files]	This is supported for transfer of restore relationship only. This specifies the list of files or LUNs to be restored. Can contain up to eight files or LUNs.
relationship	relationship	

Name	Type	Description
snapshot	string	Name of Snapshot copy being transferred.
source_snapshot	string	Specifies the Snapshot copy on the source to be transferred to the destination.
storage_efficiency_enabled	boolean	This is supported for transfer of restore relationship only. Set this property to "false" to turn off storage efficiency for data transferred over the wire and written to the destination.
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve attributes of an ongoing SnapMirror transfer

GET /snapmirror/relationships/{relationship.uuid}/transfers/{uuid}

Introduced In: 9.6

Retrieves the attributes of a specific ongoing SnapMirror transfer.

Related ONTAP commands

- `snapmirror show`

Example

```
GET "/api/snapmirror/relationships/293baa53-e63d-11e8-bff1-005056a793dd/transfers/293baa53-e63d-11e8-bff1-005056a793dd"
```

Learn more

- [DOC /snapmirror/relationships/{relationship.uuid}/transfers](#)

Parameters

Name	Type	In	Required	Description
relationship.uuid	string	path	True	Relationship UUID
uuid	string	path	True	Transfer UUID
fields	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
_links	_links	
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
relationship	relationship	
snapshot	string	Name of Snapshot copy being transferred.
state	string	Status of the transfer. Set PATCH state to "aborted" to abort the transfer. Set PATCH state to "hard_aborted" to abort the transfer and discard the restart checkpoint.

Name	Type	Description
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "relationship": {
    "destination": {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "consistency_group_volumes": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ],
      "path": "svm1:volume1",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
    },
    "snapshot": "string",
```

```
"state": "string",
"storage_efficiency_enabled": null,
"uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

files

Specifies a file or LUN consisting of a source_path and an optional destination_path. If not specified, the destination_path is the same as the source_path.

Name	Type	Description
destination_path	string	
source_path	string	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

consistency_group_volumes

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
cluster	cluster	
consistency_group_volumes	array[consistency_group_volumes]	Mandatory property for a Consistency Group endpoint. Specifies the list of FlexVol volumes for a Consistency Group.
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: ONTAP Consistency Group - svm1:/cg/cg_name <ul style="list-style-type: none"> • example: svm1:volume1 • Introduced in: 9.6
svm	svm	

relationship

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
restore	boolean	Is the relationship for restore?
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Cancel an ongoing SnapMirror transfer

PATCH /snapmirror/relationships/{relationship.uuid}/transfers/{uuid}

Introduced In: 9.6

Aborts an ongoing SnapMirror transfer. This operation is applicable on asynchronous SnapMirror relationships.

Related ONTAP commands

- `snapmirror abort`

Example

```
PATCH "/api/snapmirror/relationships/293baa53-e63d-11e8-bff1-005056a793dd/transfers/293baa53-e63d-11e8-bff1-005056a793dd"
'{"state":"aborted"}'
```

Learn more

- [DOC /snapmirror/relationships/{relationship.uuid}/transfers](#)

Parameters

Name	Type	In	Required	Description
relationship.uuid	string	path	True	Relationship UUID
uuid	string	path	True	Transfer UUID

Request Body

Name	Type	Description
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
relationship	relationship	

Name	Type	Description
snapshot	string	Name of Snapshot copy being transferred.
state	string	Status of the transfer. Set PATCH state to "aborted" to abort the transfer. Set PATCH state to "hard_aborted" to abort the transfer and discard the restart checkpoint.
uuid	string	

Example request

```
{
  "bytes_transferred": 0,
  "checkpoint_size": 0,
  "relationship": {
    "destination": {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "consistency_group_volumes": [
        {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      ],
      "path": "svm1:volume1",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      }
    },
    "uuid": "d2d7ceea-ab52-11e8-855e-00505682a4c7"
  },
  "snapshot": "string",
  "state": "string",
  "storage_efficiency_enabled": null,
  "uuid": "4ea7a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response codes

Error code	Description
13303848	Abort of sync SnapMirror is not allowed
13303849	SnapMirror transfer state is invalid

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

files

Specifies a file or LUN consisting of a source_path and an optional destination_path. If not specified, the destination_path is the same as the source_path.

Name	Type	Description
destination_path	string	
source_path	string	

cluster

Name	Type	Description
name	string	
uuid	string	

consistency_group_volumes

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapmirror_endpoint

Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.

Name	Type	Description
cluster	cluster	
consistency_group_volumes	array[consistency_group_volumes]	Mandatory property for a Consistency Group endpoint. Specifies the list of FlexVol volumes for a Consistency Group.
path	string	ONTAP FlexVol/FlexGroup - svm1:volume1 ONTAP SVM - svm1: ONTAP Consistency Group - svm1:/cg/cg_name <ul style="list-style-type: none"> • example: svm1:volume1 • Introduced in: 9.6
svm	svm	

relationship

Name	Type	Description
destination	snapmirror_endpoint	Endpoint of a SnapMirror relationship. For a GET request, the property "cluster" is populated when the endpoint is on a remote cluster. A POST request to create the destination SVM endpoint or to establish an SVM DR relationship must have the property "cluster" populated with the remote cluster details. A POST request to create the destination FlexVol volume, FlexGroup volume, and Consistency Group endpoints can optionally specify the "cluster" property when the source SVM and the destination SVM are peered. A POST request to establish a SnapMirror relationship between the source endpoint and destination endpoint and when the source SVM and the destination SVM are not peered, must specify the "cluster" property for the remote endpoint.
restore	boolean	Is the relationship for restore?
uuid	string	

snapmirror_transfer

SnapMirror transfer information

Name	Type	Description
bytes_transferred	integer	Bytes transferred
checkpoint_size	integer	Amount of data transferred in bytes as recorded in the restart checkpoint.
relationship	relationship	
snapshot	string	Name of Snapshot copy being transferred.

Name	Type	Description
state	string	Status of the transfer. Set PATCH state to "aborted" to abort the transfer. Set PATCH state to "hard_aborted" to abort the transfer and discard the restart checkpoint.
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Storage

Storage overview

Overview

The ONTAP storage APIs can be used to manage physical and logical storage. This includes management of aggregates, volumes, LUNs, qtrees, snapshots, quotas, and storage efficiency.

Retrieve or create a collection of storage aggregates

Storage aggregates endpoint overview

Retrieving storage aggregate information

The Storage Aggregate GET API retrieves all data aggregates in the cluster. System owned root aggregates

are not included in the output.

This API also supports specific queries, in addition to queries on aggregate body properties, which affect the output of the API. The parameters for these queries are "recommend" and "show_spares". Using the "recommend" query returns the list of aggregates that are recommended for creation in the cluster. The "show_spares" query returns a response outside of the records body, which includes the groups of usable spares in the cluster. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices.

The collection GET returns the aggregate identifiers, UUID and name, and the node on which the aggregate resides. The instance GET, by default, returns all of the properties defined in the aggregates object, except advanced properties.

The properties "space.footprint" and "space.block_storage.inactive_user_data" are considered advanced properties and only returned when requested using the "fields" query parameter. Performance "metric" and "statistics" for aggregates are also only returned when requested. The "statistics" property accounts for the cumulative raw values collected by ONTAP for an aggregate, while the "metric" property displays the incremental average for latency and incremental changes in IOPs and throughput over the last 15 seconds. Any external application can use the raw statistics to derive its own incremental performance metrics.

Creating storage aggregates

When the POST command is issued with no properties, the system evaluates the cluster attached storage, determines the optimal aggregate layout and configures the aggregates. This layout is completely controlled by the system.

To view the recommended optimal layout rather than creating it, use the GET endpoint, setting the "recommend" query to 'true'.

Alternatively, POST can be used with specific properties to create an aggregate as requested. At a minimum, the aggregate name, disk count, and the node where it should reside are required if any properties are provided.

When using POST with input properties, three properties are required. These are:

- name
- node.name or node.uuid
- block_storage.primary.disk_count

Remaining properties are optional

The following properties can be specified in POST:

- name - Name of the aggregate.
- node.name and node.uuid - Node on which the aggregate will be created.
- block_storage.primary.disk_count - Number of disks to be used to create the aggregate.
- block_storage.mirror.enabled - Specifies whether or not the aggregate should be created using SyncMirror.
- block_storage.primary.checksum_style - Checksum style of the disks to be use for the aggregate.
- block_storage.primary.disk_class - Class of disks to be use to for the aggregate.
- block_storage.primary.raid_size - Desired RAID size of the aggregate.
- block_storage.primary.raid_type - Desired RAID type of the aggregate.

- snaplock_type - SnapLock type to use on the aggregate.
 - data_encryption.software_encryption_enabled - Enable or disable NAE (NetApp Aggregate Encryption) on the aggregate.
-

Examples

Retrieving a list of aggregates from the cluster

The following example shows the response with a list of data aggregates in the cluster:

```
# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates" -H "accept:
application/json"

# The response:
{
  "records": [
    {
      "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
      "name": "test1",
      "node": {
        "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
        "name": "node-1",
      },
    },
    {
      "uuid": "4a7e4139-ca7a-420b-9a11-3f040d2189fd",
      "name": "test4",
      "node": {
        "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
        "name": "node-2",
      },
    },
  ],
  "num_records": 2,
}
```

Retrieving a list of aggregates recommended for creation from the cluster

The following example shows the response with a list of recommended data aggregates in the cluster.



Each aggregate UUID provided in this response is not guaranteed to be the same UUID for the aggregate if it is created.

```
# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/aggregates?recommend=true&fields=*" -H "accept:
application/json"

# The response:
{
  "records": [
    {
      "uuid": "795bf7c2-fa4b-11e8-ba65-005056bbe5c1",
      "name": "node_2_SSD_1",
      "node": {
        "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
        "name": "node-2",
      },
      "space": {
        "block_storage": {
          "size": 1116180480
        }
      },
      "block_storage": {
        "primary": {
          "disk_count": 23,
          "disk_class": "solid_state",
          "raid_type": "raid_dp",
          "disk_type": "ssd"
        },
        "hybrid_cache": {
          "enabled": false
        },
        "mirror": {
          "enabled": false
        }
      },
    },
    {
      "uuid": "795c0a15-fa4b-11e8-ba65-005056bbe5c1",
      "name": "node_1_SSD_1",
      "node": {
```

```

    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
    "name": "node-1",
  },
  "space": {
    "block_storage": {
      "size": 176238592
    }
  },
  "block_storage": {
    "primary": {
      "disk_count": 5,
      "disk_class": "solid_state",
      "raid_type": "raid_dp",
      "disk_type": "ssd"
    },
    "hybrid_cache": {
      "enabled": false
    },
    "mirror": {
      "enabled": false
    }
  },
}
],
"num_records": 2,
}

```

Retrieve a collection of aggregates for an entire cluster

GET /storage/aggregates

Introduced In: 9.6

Retrieves the collection of aggregates for the entire cluster.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `metric.*`
- `space.block_storage.inactive_user_data`
- `space.footprint`
- `statistics.*`

Related ONTAP commands

- `storage aggregate show`

Parameters

Name	Type	In	Required	Description
recommend	boolean	query	False	If set to 'true', it queries the system for the recommended optimal layout for creating new aggregates. The default setting is 'false'.
show_spares	boolean	query	False	If set to 'true', the spares object is returned instead of records to show the spare groups in the cluster. The default setting is 'false'.
metric.timestamp	string	query	False	Filter by metric.timestamp <ul style="list-style-type: none">• Introduced in: 9.7
metric.duration	string	query	False	Filter by metric.duration <ul style="list-style-type: none">• Introduced in: 9.7
metric.throughput.total	integer	query	False	Filter by metric.throughput.total <ul style="list-style-type: none">• Introduced in: 9.7
metric.throughput.read	integer	query	False	Filter by metric.throughput.read <ul style="list-style-type: none">• Introduced in: 9.7

Name	Type	In	Required	Description
metric.throughput.other	integer	query	False	Filter by metric.throughput.other • Introduced in: 9.7
metric.throughput.write	integer	query	False	Filter by metric.throughput.write • Introduced in: 9.7
metric.latency.total	integer	query	False	Filter by metric.latency.total • Introduced in: 9.7
metric.latency.read	integer	query	False	Filter by metric.latency.read • Introduced in: 9.7
metric.latency.other	integer	query	False	Filter by metric.latency.other • Introduced in: 9.7
metric.latency.write	integer	query	False	Filter by metric.latency.write • Introduced in: 9.7
metric.status	string	query	False	Filter by metric.status • Introduced in: 9.7
metric.iops.total	integer	query	False	Filter by metric.iops.total • Introduced in: 9.7

Name	Type	In	Required	Description
metric.iops.read	integer	query	False	Filter by metric.iops.read • Introduced in: 9.7
metric.iops.other	integer	query	False	Filter by metric.iops.other • Introduced in: 9.7
metric.iops.write	integer	query	False	Filter by metric.iops.write • Introduced in: 9.7
uuid	string	query	False	Filter by uuid
name	string	query	False	Filter by name
space.ency_wit hout_snapshots.logical_used	integer	query	False	Filter by space.ency_wit hout_snapshots.logical_used
space.ency_wit hout_snapshots.savings	integer	query	False	Filter by space.ency_wit hout_snapshots.savings
space.ency_wit hout_snapshots.ratio	number	query	False	Filter by space.ency_wit hout_snapshots.ratio
space.footprint	integer	query	False	Filter by space.footprint
space.block_storage.inactive_user_data	integer	query	False	Filter by space.block_storage.inactive_user_data
space.block_storage.size	integer	query	False	Filter by space.block_storage.size

Name	Type	In	Required	Description
space.block_storage.full_threshold_percent	integer	query	False	Filter by space.block_storage.full_threshold_percent
space.block_storage.available	integer	query	False	Filter by space.block_storage.available
space.block_storage.used	integer	query	False	Filter by space.block_storage.used
space.efficiency.logical_used	integer	query	False	Filter by space.efficiency.logical_used
space.efficiency.savings	integer	query	False	Filter by space.efficiency.savings
space.efficiency.ratio	number	query	False	Filter by space.efficiency.ratio
space.cloud_storage.used	integer	query	False	Filter by space.cloud_storage.used
snaplock_type	string	query	False	Filter by snaplock_type
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
block_storage.mirror.state	string	query	False	Filter by block_storage.mirror.state
block_storage.mirror.enabled	boolean	query	False	Filter by block_storage.mirror.enabled
block_storage.hybrid_cache.raid_type	string	query	False	Filter by block_storage.hybrid_cache.raid_type

Name	Type	In	Required	Description
block_storage.hybrid_cache.used	integer	query	False	Filter by block_storage.hybrid_cache.used
block_storage.hybrid_cache.disk_count	integer	query	False	Filter by block_storage.hybrid_cache.disk_count
block_storage.hybrid_cache.enabled	boolean	query	False	Filter by block_storage.hybrid_cache.enabled
block_storage.hybrid_cache.size	integer	query	False	Filter by block_storage.hybrid_cache.size
block_storage.plexes.name	string	query	False	Filter by block_storage.plexes.name
block_storage.primary.checksum_style	string	query	False	Filter by block_storage.primary.checksum_style
block_storage.primary.raid_size	integer	query	False	Filter by block_storage.primary.raid_size
block_storage.primary.disk_type	string	query	False	Filter by block_storage.primary.disk_type <ul style="list-style-type: none"> Introduced in: 9.7
block_storage.primary.disk_class	string	query	False	Filter by block_storage.primary.disk_class
block_storage.primary.disk_count	integer	query	False	Filter by block_storage.primary.disk_count
block_storage.primary.raid_type	string	query	False	Filter by block_storage.primary.raid_type

Name	Type	In	Required	Description
home_node.name	string	query	False	Filter by home_node.name
home_node.uuid	string	query	False	Filter by home_node.uuid
dr_home_node.uuid	string	query	False	Filter by dr_home_node.uuid
dr_home_node.name	string	query	False	Filter by dr_home_node.name
data_encryption.software_encryption_enabled	boolean	query	False	Filter by data_encryption.software_encryption_enabled
data_encryption.drive_protection_enabled	boolean	query	False	Filter by data_encryption.drive_protection_enabled
inactive_data_reporting.start_time	string	query	False	Filter by inactive_data_reporting.start_time • Introduced in: 9.8
inactive_data_reporting.enabled	boolean	query	False	Filter by inactive_data_reporting.enabled • Introduced in: 9.8
create_time	string	query	False	Filter by create_time
statistics.iops_raw.total	integer	query	False	Filter by statistics.iops_raw.total • Introduced in: 9.7

Name	Type	In	Required	Description
statistics.iops_raw.read	integer	query	False	Filter by statistics.iops_raw.read • Introduced in: 9.7
statistics.iops_raw.other	integer	query	False	Filter by statistics.iops_raw.other • Introduced in: 9.7
statistics.iops_raw.write	integer	query	False	Filter by statistics.iops_raw.write • Introduced in: 9.7
statistics.latency_raw.total	integer	query	False	Filter by statistics.latency_raw.total • Introduced in: 9.7
statistics.latency_raw.read	integer	query	False	Filter by statistics.latency_raw.read • Introduced in: 9.7
statistics.latency_raw.other	integer	query	False	Filter by statistics.latency_raw.other • Introduced in: 9.7
statistics.latency_raw.write	integer	query	False	Filter by statistics.latency_raw.write • Introduced in: 9.7

Name	Type	In	Required	Description
statistics.timestamp	string	query	False	Filter by statistics.timestamp • Introduced in: 9.7
statistics.status	string	query	False	Filter by statistics.status • Introduced in: 9.7
statistics.throughput_raw.total	integer	query	False	Filter by statistics.throughput_raw.total • Introduced in: 9.7
statistics.throughput_raw.read	integer	query	False	Filter by statistics.throughput_raw.read • Introduced in: 9.7
statistics.throughput_raw.other	integer	query	False	Filter by statistics.throughput_raw.other • Introduced in: 9.7
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write • Introduced in: 9.7
state	string	query	False	Filter by state
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
error	error	
num_records	integer	Number of records
records	array[aggregate]	
spares	array[aggregate_spare]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "block_storage": {
        "hybrid_cache": {
          "disk_count": "6",
          "raid_type": "string",
          "size": "1612709888",
          "used": "26501122"
        },
        "mirror": {
          "enabled": "",
          "state": "string"
        },
        "plexes": [
          {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            }
          }
        ]
      }
    }
  ]
}
```

```

    },
    "name": "plex0"
  }
],
"primary": {
  "checksum_style": "string",
  "disk_class": "performance",
  "disk_count": "8",
  "disk_type": "string",
  "raid_size": "16",
  "raid_type": "string"
}
},
"cloud_storage": {
  "stores": [
    {
      "cloud_store": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "store1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "used": 0
    }
  ],
  "create_time": "2018-01-01T12:00:00-04:00",
  "dr_home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "home_node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "inactive_data_reporting": {
    "start_time": "2019-12-12T12:00:00-04:00"
  },

```



```

"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"snaplock_type": "string",
"space": {
  "block_storage": {
    "available": "10156560384",
    "full_threshold_percent": 0,
    "inactive_user_data": "304448",
    "size": "10156769280",
    "used": "2088960"
  },
  "cloud_storage": {
    "used": "402743264"
  }
}

```

```

    },
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "footprint": "608896"
  },
  "state": "string",
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "uuid": "string"
}
],
"spares": [
  {
    "checksum_style": "string",
    "disk_class": "solid_state",
    "layout_requirements": [
      {
        "aggregate_min_disks": "6",
        "raid_group": {
          "default": "16",
          "max": "28",

```

```

        "min": "5"
      },
      "raid_type": "string"
    }
  ],
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": "10156769280",
  "syncmirror_pool": "string",
  "usable": "9"
}
]
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
787092	The target field cannot be specified for this operation.
8586225	Encountered unexpected error in retrieving metrics and statistics for an aggregate.
19726341	Not enough eligible spare disks are available on the node.
19726344	No recommendation can be made for this cluster.
19726357	Aggregate recommendations are not supported on MetroCluster.
19726358	Aggregate recommendations are not supported on ONTAP Cloud.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.

Error Code	Description
19726386	Encountered an error when retrieving licensing information on this cluster.
19726387	No recommendation can be provided for this cluster within the license capacity.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

_links

Name	Type	Description
self	href	

hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
enabled	boolean	Specifies whether the aggregate uses HDDs with SSDs as a cache.
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
checksum_style	string	The checksum style used by the aggregate.

Name	Type	Description
disk_class	string	The class of disks being used by the aggregate.
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
disk_type	string	The type of disk being used by the aggregate.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.

block_storage

Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
hybrid_cache	hybrid_cache	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	mirror	
plexes	array[plex_reference]	Plex reference for each plex in the aggregate.
primary	primary	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	

Name	Type	Description
uuid	string	

cloud_storage_tier

Name	Type	Description
cloud_store	cloud_store	Cloud store
used	integer	Capacity used in bytes in the cloud store by this aggregate. This is a cached value calculated every 5 minutes.

cloud_storage

Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.

Name	Type	Description
attach_eligible	boolean	Specifies whether the aggregate is eligible for a cloud store to be attached.
stores	array[cloud_storage_tier]	Configuration information for each cloud storage portion of the aggregate.
tiering_fullness_threshold	integer	The percentage of space in the performance tier that must be used before data is tiered out to the cloud store. Only valid for PATCH operations.

data_encryption

Name	Type	Description
drive_protection_enabled	boolean	Specifies whether the aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	Specifies whether NetApp aggregate encryption is enabled. All data in the aggregate is encrypted.

dr_home_node

Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.

Name	Type	Description
name	string	
uuid	string	

home_node

Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

inactive_data_reporting

Name	Type	Description
enabled	boolean	Specifies whether or not inactive data reporting is enabled on the aggregate.
start_time	string	Timestamp at which inactive data reporting was enabled on the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

The most recent sample of I/O metrics for the aggregate.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

block_storage

Name	Type	Description
available	integer	Space available in bytes.
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data</code> or <code>**</code> .
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for aggregates with a cloud store tier.

efficiency

Storage efficiency

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

space

Name	Type	Description
block_storage	block_storage	
cloud_storage	cloud_storage	
efficiency	efficiency	Storage efficiency
efficiency_without_snapshots	efficiency_without_snapshots	Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
footprint	integer	A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or **.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

The real time I/O statistics for the aggregate.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

aggregate

Name	Type	Description
_links	_links	
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
metric	metric	The most recent sample of I/O metrics for the aggregate.
name	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
statistics	statistics	The real time I/O statistics for the aggregate.

Name	Type	Description
uuid	string	Aggregate UUID

raid_group

Name	Type	Description
default	integer	Default number of disks in a RAID group
max	integer	Maximum number of disks allowed in a RAID group
min	integer	Minimum number of disks allowed in a RAID group

layout_requirement

Name	Type	Description
aggregate_min_disks	integer	Minimum number of disks to create an aggregate
default	boolean	Indicates if this RAID type is the default
raid_group	raid_group	
raid_type	string	RAID type

node

Node where the spares are assigned

Name	Type	Description
_links	_links	
name	string	
uuid	string	

aggregate_spare

Name	Type	Description
checksum_style	string	The checksum type that has been assigned to the spares

Name	Type	Description
disk_class	string	Disk class of spares
layout_requirements	array[layout_requirement]	Available RAID protections and their restrictions
node	node	Node where the spares are assigned
size	integer	Usable size of each spare in bytes
syncmirror_pool	string	SyncMirror spare pool
usable	integer	<p>Total number of usable spares in the bucket. The usable count for each class of spares does not include reserved spare capacity recommended by ONTAP best practices.</p> <ul style="list-style-type: none"> • example: 9 • readOnly: 1 • Introduced in: 9.6

Create a collection of aggregates for an entire cluster

POST /storage/aggregates

Introduced In: 9.6

Automatically creates aggregates based on an optimal layout recommended by the system. Alternatively, properties can be provided to create an aggregate according to the requested specification. This request starts a job and returns a link to that job.

POST operations will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation.

Required properties

Properties are not required for this API. The following properties are only required if you want to specify properties for aggregate creation:

- name - Name of the aggregate.
- node.name or node.uuid - Node on which the aggregate will be created.
- block_storage.primary.disk_count - Number of disks to be used to create the aggregate.

Default values

If not specified in POST, the following default values are assigned. The remaining unspecified properties will receive system dependent default values.

- `block_storage.mirror.enabled` - *false*
- `snaplock_type` - *non_snaplock*

Related ONTAP commands

- `storage aggregate auto-provision`
- `storage aggregate create`

Example:

```
POST /api/storage/aggregates {"node": {"name": "node1"}, "name": "test",  
"block_storage": {"primary": {"disk_count": "10"}}
```

Parameters

Name	Type	In	Required	Description
disk_size	integer	query	False	If set, POST only selects disks of the specified size.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
name	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

Example request

```
{
  "block_storage": {
    "hybrid_cache": {
      "disk_count": "6",
      "raid_type": "string",
      "size": "1612709888",
      "used": "26501122"
    },
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "name": "plex0"
      }
    ],
    "primary": {
      "checksum_style": "string",
      "disk_class": "performance",
      "disk_count": "8",
      "disk_type": "string",
      "raid_size": "16",
      "raid_type": "string"
    }
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01T12:00:00-04:00",
  "dr_home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "home_node": {
    "name": "node1",
```

```

    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "inactive_data_reporting": {
    "enabled": null,
    "start_time": "2019-12-12T12:00:00-04:00"
  },
  "name": "node1_aggr_1",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snaplock_type": "string",
  "space": {
    "block_storage": {
      "available": "10156560384",
      "full_threshold_percent": 0,
      "inactive_user_data": "304448",
      "size": "10156769280",
      "used": "2088960"
    },
    "cloud_storage": {
      "used": "402743264"
    },
    "efficiency": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "efficiency_without_snapshots": {
      "logical_used": 0,
      "ratio": 0,
      "savings": 0
    },
    "footprint": "608896"
  },
  "state": "string",
  "uuid": "string"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
460770	The aggregate create job failed to create the aggregate.
786438	Failed to create an aggregate on the node.
786439	An aggregate already uses the specified name.
786446	The node is not in cluster.
786468	VLDB is offline.
786819	The value is invalid for the specified option at the current privilege level.
786902	RAID-TEC aggregate is not fully supported.
786911	Not every node in the cluster has the Data ONTAP version required for the feature.
787069	Node is setup for MetroCluster over IP configuration; creating an unmirrored aggregate is not supported in this configuration.
787092	The target field cannot be specified for this operation.
1114292	The required SnapLock license is not installed.
2425736	No matching node found for the target UUID.
19726341	Not enough eligible spare disks are available on the node.
19726344	No recommendation can be made for this cluster.

Error Code	Description
19726357	Automatic aggregate creation is not supported on MetroCluster.
19726358	Automatic aggregate creation is not supported on ONTAP Cloud.
19726373	Recommendation specified for creating aggregates is not current.
19726378	Failed to create recommended aggregates on one or more nodes.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
19726386	Encountered an error when retrieving licensing information on this cluster.
19726387	No recommendation can be provided for this cluster within the license capacity.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
enabled	boolean	Specifies whether the aggregate uses HDDs with SSDs as a cache.
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
name	string	

primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
checksum_style	string	The checksum style used by the aggregate.
disk_class	string	The class of disks being used by the aggregate.
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
disk_type	string	The type of disk being used by the aggregate.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.

block_storage

Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
hybrid_cache	hybrid_cache	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	mirror	
plexes	array[plex_reference]	Plex reference for each plex in the aggregate.

Name	Type	Description
primary	primary	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud_store

Cloud store

Name	Type	Description
name	string	
uuid	string	

cloud_storage_tier

Name	Type	Description
cloud_store	cloud_store	Cloud store
used	integer	Capacity used in bytes in the cloud store by this aggregate. This is a cached value calculated every 5 minutes.

cloud_storage

Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.

Name	Type	Description
attach_eligible	boolean	Specifies whether the aggregate is eligible for a cloud store to be attached.
stores	array[cloud_storage_tier]	Configuration information for each cloud storage portion of the aggregate.
tiering_fullness_threshold	integer	The percentage of space in the performance tier that must be used before data is tiered out to the cloud store. Only valid for PATCH operations.

data_encryption

Name	Type	Description
drive_protection_enabled	boolean	Specifies whether the aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	Specifies whether NetApp aggregate encryption is enabled. All data in the aggregate is encrypted.

dr_home_node

Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.

Name	Type	Description
name	string	
uuid	string	

home_node

Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
name	string	
uuid	string	

inactive_data_reporting

Name	Type	Description
start_time	string	Timestamp at which inactive data reporting was enabled on the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Performance metric for write I/O operations.

metric

The most recent sample of I/O metrics for the aggregate.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node

Node where the aggregate currently resides.

Name	Type	Description
name	string	
uuid	string	

block_storage

Name	Type	Description
available	integer	Space available in bytes.
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data</code> or <code>**</code> .
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.

Name	Type	Description
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for aggregates with a cloud store tier.

efficiency

Storage efficiency

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

space

Name	Type	Description
block_storage	block_storage	
cloud_storage	cloud_storage	

Name	Type	Description
efficiency	efficiency	Storage efficiency
efficiency_without_snapshots	efficiency_without_snapshots	Storage efficiency that does not include the savings provided by Snapshot copies.
footprint	integer	A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or **.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

The real time I/O statistics for the aggregate.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

aggregate

Name	Type	Description
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
name	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a collection of cloud stores used by an aggregate

GET /storage/aggregates/{aggregate.uuid}/cloud-stores

Introduced In: 9.6

Retrieves the collection of cloud stores used by an aggregate.

Related ONTAP commands

- `storage aggregate object-store show`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
target.name	string	query	False	Filter by target.name

Name	Type	In	Required	Description
target.uuid	string	query	False	Filter by target.uuid
unavailable_reason. message	string	query	False	Filter by unavailable_reason. message <ul style="list-style-type: none"> Introduced in: 9.7
used	integer	query	False	Filter by used
unreclaimed_space_ threshold	integer	query	False	Filter by unreclaimed_space_ threshold
mirror_degraded	boolean	query	False	Filter by mirror_degraded
availability	string	query	False	Filter by availability
primary	boolean	query	False	Filter by primary
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[cloud_store]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "availability": "string",
      "target": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "target1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "unavailable_reason": {
        "message": "string"
      },
      "unreclaimed_space_threshold": "20",
      "used": 0
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

unavailable_reason

Name	Type	Description
message	string	Indicates why the object store is unavailable.

cloud_store

Name	Type	Description
_links	_links	
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.

Name	Type	Description
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Attach an object store to an aggregate or add a second object store as a mirror

POST /storage/aggregates/{aggregate.uuid}/cloud-stores

Introduced In: 9.6

Attaches an object store to an aggregate, or adds a second object store as a mirror.

Required properties

- `target.uuid` or `target.name` - UUID or name of the cloud target.

Recommended optional properties

- `primary` - *true* if the object store is primary or *false* if it is a mirror.
- `allow_flexgroups` - Allow attaching object store to an aggregate containing FlexGroup constituents.
- `check_only` - Validate only and do not add the cloud store.

Default property values

- `primary` - *true*
- `allow_flexgroups` - *false*
- `check_only` - *false*

Related ONTAP commands

- `storage aggregate object-store attach`
- `storage aggregate object-store mirror`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
allow_flexgroups	boolean	query	False	This optional parameter allows attaching object store to an aggregate containing FlexGroup constituents. The default value is false. Mixing FabricPools and non-FabricPools within a FlexGroup is not recommended. All aggregates hosting constituents of a FlexGroup should be attached to the object store.

Name	Type	In	Required	Description
check_only	boolean	query	False	Validate only and do not add the cloud store.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
availability	string	Availability of the object store.

Name	Type	Description
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

Example request

```
{
  "availability": "string",
  "target": {
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": "20",
  "used": 0
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

target

Cloud target

Name	Type	Description
name	string	
uuid	string	

unavailable_reason

Name	Type	Description
message	string	Indicates why the object store is unavailable.

cloud_store

Name	Type	Description
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.

Name	Type	Description
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Remove a cloud target from an aggregate

```
DELETE /storage/aggregates/{aggregate.uuid}/cloud-stores/{target.uuid}
```

Introduced In: 9.6

Removes the specified cloud target from the aggregate. Only removal of a mirror is allowed. The primary cannot be removed. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate object-store unmirror`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
target.uuid	string	path	True	Cloud target UUID
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve the cloud store for an aggregate

GET /storage/aggregates/{aggregate.uuid}/cloud-stores/{target.uuid}

Introduced In: 9.6

Retrieves the cloud store for the aggregate using the specified cloud target UUID.

Related ONTAP commands

- `storage aggregate object-store show`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
target.uuid	string	path	True	Cloud target UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
target	target	Cloud target
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.

Name	Type	Description
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "availability": "string",
  "target": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "target1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": "20",
  "used": 0
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

target

Cloud target

Name	Type	Description
_links	_links	
name	string	
uuid	string	

unavailable_reason

Name	Type	Description
message	string	Indicates why the object store is unavailable.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Update a cloud store

PATCH /storage/aggregates/{aggregate.uuid}/cloud-stores/{target.uuid}

Introduced In: 9.6

Updates the cloud store specified by the UUID with the fields in the body. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate object-store modify`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
target.uuid	string	path	True	Cloud target UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
unavailable_reason	unavailable_reason	

Name	Type	Description
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

Example request

```
{
  "availability": "string",
  "unavailable_reason": {
    "message": "string"
  },
  "unreclaimed_space_threshold": "20",
  "used": 0
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

target

Cloud target

Name	Type	Description
name	string	
uuid	string	

unavailable_reason

Name	Type	Description
message	string	Indicates why the object store is unavailable.

cloud_store

Name	Type	Description
availability	string	Availability of the object store.
mirror_degraded	boolean	This field identifies if the mirror cloud store is in sync with the primary cloud store of a FabricPool.
primary	boolean	This field indicates whether the cloud store is the primary cloud store of a mirrored FabricPool.
unavailable_reason	unavailable_reason	
unreclaimed_space_threshold	integer	Usage threshold for reclaiming unused space in the cloud store. Valid values are 0 to 99. The default value depends on the provider type. This can be specified in PATCH but not POST.

Name	Type	Description
used	integer	The amount of object space used. Calculated every 5 minutes and cached.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage storage aggregate plexes

Storage aggregates aggregate.uuid plexes endpoint overview

Overview

The Storage Aggregate Plex API provides relevant state information for each plex in the aggregate. For each plex, details are provided for the RAID groups in the plex and the disks that make up each RAID group.

Examples

Retrieving the list of plexes in an aggregate

The following example shows the response with the list of plexes in an aggregate:

```
# The API:
/api/storage/aggregates/{uuid}/plexes

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c/plexes" -H "accept: application/json"

# The response:
{
  "records": [
    {
      "name": "plex0",
    },
    {
      "name": "plex4",
    }
  ],
  "num_records": 2,
}
```

Retrieving a specific plex in an aggregate

The following example shows the response when requesting a specific plex of an aggregate:

```
# The API:
/api/storage/aggregates/{uuid}/plexes/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c/plexes/plex0" -H "accept: application/json"

# The response:
{
  "aggregate": {
    "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
    "name": "test1",
  },
  "name": "plex0",
  "online": true,
  "state": "normal",
  "pool": "pool0",
  "resync": {
```

```
"active": false
},
"raid_groups": [
  {
    "name": "rg0",
    "cache_tier": false,
    "degraded": false,
    "recomputing_parity": {
      "active": false
    },
    "reconstruct": {
      "active": false
    },
    "disks": [
      {
        "position": "dparity",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.29",
        }
      },
      {
        "position": "parity",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.4",
        }
      },
      {
        "position": "data",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
        "disk": {
          "name": "1.1.30",
        }
      },
      {
        "position": "data",
        "state": "normal",
        "type": "ssd",
        "usable_size": 86769664,
```

```

    "disk": {
      "name": "1.1.5",
    }
  },
  {
    "position": "data",
    "state": "normal",
    "type": "ssd",
    "usable_size": 86769664,
    "disk": {
      "name": "1.1.31",
    }
  },
  {
    "position": "data",
    "state": "normal",
    "type": "ssd",
    "usable_size": 86769664,
    "disk": {
      "name": "1.1.6",
    }
  }
]
}
],
}

```

Retrieve a collection of plexes for an aggregate

GET /storage/aggregates/{aggregate.uuid}/plexes

Introduced In: 9.6

Retrieves the collection of plexes for the specified aggregate.

Related ONTAP commands

- storage aggregate plex show

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
online	boolean	query	False	Filter by online

Name	Type	In	Required	Description
name	string	query	False	Filter by name
pool	string	query	False	Filter by pool
resync.level	string	query	False	Filter by resync.level
resync.active	boolean	query	False	Filter by resync.active
resync.percent	integer	query	False	Filter by resync.percent
aggregate.uuid	string	query	False	Filter by aggregate.uuid
aggregate.name	string	query	False	Filter by aggregate.name
raid_groups.reconstruct.percent	integer	query	False	Filter by raid_groups.reconstruct.percent • Introduced in: 9.7
raid_groups.reconstruct.active	boolean	query	False	Filter by raid_groups.reconstruct.active • Introduced in: 9.7
raid_groups.recomputing_parity.active	boolean	query	False	Filter by raid_groups.recomputing_parity.active • Introduced in: 9.7
raid_groups.recomputing_parity.percent	integer	query	False	Filter by raid_groups.recomputing_parity.percent • Introduced in: 9.7

Name	Type	In	Required	Description
raid_groups.degraded	boolean	query	False	Filter by raid_groups.degraded • Introduced in: 9.7
raid_groups.name	string	query	False	Filter by raid_groups.name • Introduced in: 9.7
raid_groups.disks.usable_size	integer	query	False	Filter by raid_groups.disks.usable_size • Introduced in: 9.7
raid_groups.disks.disk.name	string	query	False	Filter by raid_groups.disks.disk.name • Introduced in: 9.7
raid_groups.disks.state	string	query	False	Filter by raid_groups.disks.state • Introduced in: 9.7
raid_groups.disks.type	string	query	False	Filter by raid_groups.disks.type • Introduced in: 9.7
raid_groups.disks.position	string	query	False	Filter by raid_groups.disks.position • Introduced in: 9.7

Name	Type	In	Required	Description
raid_groups.cache_tier	boolean	query	False	Filter by raid_groups.cache_tier • Introduced in: 9.7
state	string	query	False	Filter by state
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
error	error	
num_records	integer	Number of records
records	array[plex]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "records": [
    {
      "aggregate": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-aelc-123478563412"
      },
      "name": "plex0",
      "pool": "string",
      "raid_groups": [
        {
          "disks": [
            {
              "disk": {
                "_links": {
                  "self": {
                    "href": "/api/resourcelink"
                  }
                }
              },
              "name": "1.0.1"
            }
          ]
        }
      ]
    }
  ]
}
```



```

        },
        "position": "string",
        "state": "string",
        "type": "ssd",
        "usable_size": "947912704"
    }
],
"name": "rg0",
"recomputing_parity": {
    "percent": "10"
},
"reconstruct": {
    "percent": "10"
}
}
],
"resync": {
    "level": "string",
    "percent": "10"
},
"state": "string"
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

_links

Name	Type	Description
self	href	

aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	

Name	Type	Description
uuid	string	

disk

Disk

Name	Type	Description
_links	_links	
name	string	

raid_group_disk

Name	Type	Description
disk	disk	Disk
position	string	The position of the disk within the RAID group.
state	string	The state of the disk within the RAID group.
type	string	Disk interface type
usable_size	integer	Size in bytes that is usable by the aggregate.

recomputing_parity

Name	Type	Description
active	boolean	RAID group is recomputing parity
percent	integer	Recomputing parity percentage

reconstruct

Name	Type	Description
active	boolean	One or more disks in this RAID group are being reconstructed.
percent	integer	Reconstruct percentage

raid_group

Name	Type	Description
cache_tier	boolean	RAID group is a cache tier
degraded	boolean	RAID group is degraded. A RAID group is degraded when at least one disk from that group has failed or is offline.
disks	array[raid_group_disk]	
name	string	RAID group name
recomputing_parity	recomputing_parity	
reconstruct	reconstruct	

resync

Name	Type	Description
active	boolean	Plex is being resynchronized to its mirrored plex
level	string	Plex resyncing level
percent	integer	Plex resyncing percentage

plex

Name	Type	Description
aggregate	aggregate	Aggregate
name	string	Plex name
online	boolean	Plex is online
pool	string	SyncMirror pool assignment
raid_groups	array[raid_group]	
resync	resync	
state	string	Plex state

Retrieve a plex specified by the aggregate UUID and plex name

GET /storage/aggregates/{aggregate.uuid}/plexes/{name}

Introduced In: 9.6

Retrieves the plex specified by the aggregate UUID and plex name.

Related ONTAP commands

- `storage aggregate plex show`

Parameters

Name	Type	In	Required	Description
aggregate.uuid	string	path	True	Aggregate UUID
name	string	path	True	Plex name
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
aggregate	aggregate	Aggregate
name	string	Plex name
online	boolean	Plex is online
pool	string	SyncMirror pool assignment
raid_groups	array[raid_group]	
resync	resync	
state	string	Plex state

Example response

```

{
  "aggregate": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "name": "plex0",
  "pool": "string",
  "raid_groups": [
    {
      "disks": [
        {
          "disk": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "1.0.1"
          },
          "position": "string",
          "state": "string",
          "type": "ssd",
          "usable_size": "947912704"
        }
      ],
      "name": "rg0",
      "recomputing_parity": {
        "percent": "10"
      },
      "reconstruct": {
        "percent": "10"
      }
    }
  ],
  "resync": {
    "level": "string",
    "percent": "10"
  },
  "state": "string"
}

```


Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

disk

Disk

Name	Type	Description
_links	_links	
name	string	

raid_group_disk

Name	Type	Description
disk	disk	Disk
position	string	The position of the disk within the RAID group.
state	string	The state of the disk within the RAID group.
type	string	Disk interface type
usable_size	integer	Size in bytes that is usable by the aggregate.

recomputing_parity

Name	Type	Description
active	boolean	RAID group is recomputing parity
percent	integer	Recomputing parity percentage

reconstruct

Name	Type	Description
active	boolean	One or more disks in this RAID group are being reconstructed.
percent	integer	Reconstruct percentage

raid_group

Name	Type	Description
cache_tier	boolean	RAID group is a cache tier
degraded	boolean	RAID group is degraded. A RAID group is degraded when at least one disk from that group has failed or is offline.
disks	array[raid_group_disk]	
name	string	RAID group name
recomputing_parity	recomputing_parity	
reconstruct	reconstruct	

resync

Name	Type	Description
active	boolean	Plex is being resynchronized to its mirrored plex
level	string	Plex resyncing level
percent	integer	Plex resyncing percentage

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage storage aggregates

Storage aggregates UUID endpoint overview

Updating storage aggregates

The PATCH operation is used to modify properties of the aggregate. There are several properties that can be modified on an aggregate. Only one property can be modified for each PATCH request.

PATCH operations on the aggregate's disk count will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation.

The following is a list of properties that can be modified using the PATCH operation including a brief description for each:

- name - This property can be changed to rename the aggregate.
- node.name and node.uuid - Either property can be updated in order to relocate the aggregate to a different node in the cluster.
- block_storage.mirror.enabled - This property can be changed from 'false' to 'true' in order to mirror the aggregate, if the system is capable of doing so.
- block_storage.primary.disk_count - This property can be updated to increase the number of disks in an aggregate.
- block_storage.primary.raid_size - This property can be updated to set the desired RAID size.
- block_storage.primary.raid_type - This property can be updated to set the desired RAID type.
- cloud_storage.tiering_fullness_threshold - This property can be updated to set the desired tiering fullness threshold if using FabricPool.
- data_encryption.software_encryption_enabled - This property enables or disables NAE on the aggregate.

Aggregate expansion

The PATCH operation also supports automatically expanding an aggregate based on the spare disks which are present within the system. Running PATCH with the query "auto_provision_policy" set to "expand" starts the recommended expansion job. In order to see the expected change in capacity before starting the job, call GET on an aggregate instance with the query "auto_provision_policy" set to "expand".

Manual simulated aggregate expansion

The PATCH operation also supports simulated manual expansion of an aggregate. Running PATCH with the query "simulate" set to "true" and "block_storage.primary.disk_count" set to the final disk count will start running the prechecks associated with expanding the aggregate to the proposed size. The response body will include information on how many disks the aggregate can be expanded to, any associated warnings, along with the proposed final size of the aggregate.

Deleting storage aggregates

If volumes exist on an aggregate, they must be deleted or moved before the aggregate can be deleted. See the /storage/volumes API for details on moving or deleting volumes.

Examples

Retrieving a specific aggregate from the cluster

The following example shows the response of the requested aggregate. If there is no aggregate with the requested UUID, an error is returned.

```
# The API:
/api/storage/aggregates/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/870dd9f2-bdfa-4167-b692-57dlcec874d4" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "test1",
  "node": {
    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
    "name": "node-1",
  },
  "home_node": {
    "uuid": "caf95bec-f801-11e8-8af9-005056bbe5c1",
    "name": "node-1",
  },
  "space": {
    "block_storage": {
```

```

    "size": 235003904,
    "available": 191942656,
    "used": 43061248,
    "full_threshold_percent": 98
  },
  "cloud_storage": {
    "used": 0
  },
  "efficiency": {
    "savings": 1408029,
    "ratio": 6.908119720880661,
    "logical_used": 1646350
  },
  "efficiency_without_snapshots": {
    "savings": 0,
    "ratio": 1,
    "logical_used": 737280
  }
},
"state": "online",
"snaplock_type": "non_snaplock",
"create_time": "2018-12-04T15:40:38-05:00",
"data_encryption": {
  "software_encryption_enabled": false,
  "drive_protection_enabled": false
},
"block_storage": {
  "primary": {
    "disk_count": 6,
    "disk_class": "solid_state",
    "raid_type": "raid_dp",
    "raid_size": 24,
    "checksum_style": "block",
    "disk_type": "ssd"
  },
  "hybrid_cache": {
    "enabled": false
  },
  "mirror": {
    "enabled": false,
    "state": "unmirrored"
  },
  "plexes": [
    {
      "name": "plex0",
    }
  ]
}

```

```

    ],
    "cloud_storage": {
        "attach_eligible": false
    },
}

```

Retrieving statistics and metric for an aggregate

In this example, the API returns the "statistics" and "metric" properties for the aggregate requested.

```

#The API:
/api/storage/aggregates/{uuid}?fields=statistics,metric

#The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-bad0-005056b48388?fields=statistics,metric" -H "accept: application/json"

#The response:
{
  "uuid": "538bf337-1b2c-11e8-bad0-005056b48388",
  "name": "aggr4",
  "metric": {
    "timestamp": "2019-07-08T22:16:45Z",
    "duration": "PT15S",
    "status": "ok",
    "throughput": {
      "read": 7099,
      "write": 840226,
      "other": 193293789,
      "total": 194141115
    },
    "latency": {
      "read": 149,
      "write": 230,
      "other": 123,
      "total": 124
    },
    "iops": {
      "read": 1,
      "write": 17,
      "other": 11663,
      "total": 11682
    },
  },
  "statistics": {

```

```

    "timestamp": "2019-07-08T22:17:09Z",
    "status": "ok",
    "throughput_raw": {
      "read": 3106045952,
      "write": 63771742208,
      "other": 146185560064,
      "total": 213063348224
    },
    "latency_raw": {
      "read": 54072313,
      "write": 313354426,
      "other": 477201985,
      "total": 844628724
    },
    "iops_raw": {
      "read": 328267,
      "write": 1137230,
      "other": 1586535,
      "total": 3052032
    }
  },
}

```

For more information and examples on viewing historical performance metrics for any given aggregate, see [DOC /storage/aggregates/{uuid}/metrics](#)

Simulating aggregate expansion

The following example shows the response for a simulated data aggregate expansion based on the values of the 'block_storage.primary.disk_count' attribute passed in.

The query does not modify the existing aggregate but returns how the aggregate will look after the expansion along with any associated warnings.

Simulated data aggregate expansion will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation.

This will be reflected in the following attributes:

- space.block_storage.size - Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
- block_storage.primary.disk_count - Number of disks that could be used to create the aggregate.

```

# The API:
/api/storage/aggregates/{uuid}?simulate=true

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314?simulate=true" -H "accept: application/json" -d '{"block_storage": {"primary": {"disk_count": 14}}}'

```



```

# The response:
{
  "warnings": [
    {
      "name": "node_2_SSD_1",
      "warning": {
        "message": "One or more disks will not be added. 14 disks specified,
13 disks will be added.",
        "code": 787170,
        "arguments": [
          "14",
          "13"
        ]
      }
    }
  ],
  "num_records": 1,
  "records": [
    {
      "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
      "name": "node_2_SSD_1",
      "node": {
        "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
        "name": "node-2",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/4046dda8-f802-11e8-8f6d-
005056bb2030"
          }
        }
      },
      "space": {
        "block_storage": {
          "size": 1116180480
        }
      },
      "block_storage": {
        "primary": {
          "disk_count": 23,
          "disk_class": "solid_state",
          "raid_type": "raid_dp",
          "disk_type": "ssd"
        },
        "hybrid_cache": {
          "enabled": false
        }
      }
    }
  ]
}

```

```

    },
    "mirror": {
      "enabled": false
    }
  },
  ],
  "job": {
    "_links": {
      "self": {
        "href": "/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314"
      }
    }
  }
}

```

Retrieving a recommendation for an aggregate expansion

The following example shows the response with the recommended data aggregate expansion based on what disks are present within the system.

The query does not modify the existing aggregate but returns how the aggregate will look after the expansion. The recommendation will be reflected in the attributes - 'space.block_storage.size' and 'block_storage.primary.disk_count'.

Recommended data aggregate expansion will be blocked while one or more nodes in the cluster are simulating or implementing automatic aggregate creation.

```

# The API:
/api/storage/aggregates/{uuid}?auto_provision_policy=expand

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314?auto_provision_policy=expand" -H "accept: application/json"

# The response:
{
  "uuid": "cae60cfe-deae-42bd-babb-ef437d118314",
  "name": "node_2_SSD_1",
  "node": {
    "uuid": "4046dda8-f802-11e8-8f6d-005056bb2030",
    "name": "node-2",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/4046dda8-f802-11e8-8f6d-005056bb2030"
      }
    }
  }
}

```

```

},
"space": {
  "block_storage": {
    "size": 1116180480
  }
},
"block_storage": {
  "primary": {
    "disk_count": 23,
    "disk_class": "solid_state",
    "raid_type": "raid_dp",
    "disk_type": "ssd"
  },
  "hybrid_cache": {
    "enabled": false
  },
  "mirror": {
    "enabled": false
  }
},
"_links": {
  "self": {
    "href": "/api/storage/aggregates/cae60cfe-deae-42bd-babb-ef437d118314"
  }
}
}

```

Updating an aggregate in the cluster

The following example shows the workflow of adding disks to the aggregate.

Step 1: Check the current disk count on the aggregate.

```
# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c?fields=block_storage.primary.disk_count" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "test1",
  "block_storage": {
    "primary": {
      "disk_count": 6
    }
  },
}
```

Step 2: Update the aggregate with the new disk count in 'block_storage.primary.disk_count'. The response to PATCH is a job unless the request is invalid.

```
# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c" -H "accept: application/hal+json" -d '{"block_storage": {"primary": {"disk_count": 8}}}'

# The response:
{
  "job": {
    "uuid": "c103d15e-730b-11e8-a57f-005056b465d6",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/c103d15e-730b-11e8-a57f-005056b465d6"
      }
    }
  }
}
```

Step 3: Wait for the job to finish, then call GET to see the reflected change.

```
# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/19425837-f2fa-4a9f-8f01-712f626c983c?fields=block_storage.primary.disk_count" -H "accept: application/json"

# The response:
{
  "uuid": "19425837-f2fa-4a9f-8f01-712f626c983c",
  "name": "test1",
  "block_storage": {
    "primary": {
      "disk_count": 8
    }
  },
}
```

The following example shows the workflow to enable software encryption on an aggregate.

Step 1: Check the current software encryption status of the aggregate.

```
# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-4d93-9590-5a402bffb4b?fields=data_encryption.software_encryption_enabled" -H "accept: application/json"

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bffb4b",
  "name": "aggr5",
  "data_encryption": {
    "software_encryption_enabled": false
  },
}
```

Step 2: Update the aggregate with the encryption status in 'data_encryption.software_encryption_enabled'. The response to PATCH is a job unless the request is invalid.

```
# The API:
/api/storage/aggregates

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-4d93-9590-5a402bffb4b" -H "accept: application/hal+json" -d
'{"data_encryption": {"software_encryption_enabled": "true"}}'

# The response:
{
  "job": {
    "uuid": "6b7ab28e-168d-11ea-8a50-0050568eca76",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6b7ab28e-168d-11ea-8a50-0050568eca76"
      }
    }
  }
}
```

Step 3: Wait for the job to finish, then call GET to see the reflected change.

```
# The API:
/api/storage/aggregates

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/f3aafdc6-be35-4d93-9590-5a402bffb4b?fields=data_encryption.software_encryption_enabled" -H
"accept: application/json"

# The response:
{
  "uuid": "f3aafdc6-be35-4d93-9590-5a402bffb4b",
  "name": "aggr5",
  "data_encryption": {
    "software_encryption_enabled": true
  },
}
```

Delete an aggregate specified by the UUID

DELETE /storage/aggregates/{uuid}

Introduced In: 9.6

Deletes the aggregate specified by the UUID. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate delete`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Aggregate UUID
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
460770	The aggregate delete job failed to delete the aggregate.
460777	Failed to get information on the delete job.
786435	Internal Error. Failed to create a communication handle.
786451	Failed to delete specified aggregate.
786468	VLDB is offline.
786472	Node that hosts the aggregate is offline.
786497	Cannot delete an aggregate that has volumes.
786771	Aggregate does not exist.
786867	Specified aggregate resides on the remote cluster.
786897	Specified aggregate cannot be deleted as it is a switched-over root aggregate.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve an aggregate specified by the UUID

GET /storage/aggregates/{uuid}

Introduced In: 9.6

Retrieves the aggregate specified by the UUID. The recommend query cannot be used for this operation.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `metric.*`
- `space.block_storage.inactive_user_data`
- `space.footprint`
- `statistics.*`

Related ONTAP commands

- `storage aggregate show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Aggregate UUID
auto_provision_policy	string	query	False	If set to expand, a query is run on the system for the recommended optimal expansion layout of the aggregate. • Introduced in: 9.8
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
<code>_links</code>	_links	

Name	Type	Description
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
metric	metric	The most recent sample of I/O metrics for the aggregate.
name	string	Aggregate name
node	node	Node where the aggregate currently resides.
snaplock_type	string	SnapLock type
space	space	
state	string	Operational state of the aggregate
statistics	statistics	The real time I/O statistics for the aggregate.
uuid	string	Aggregate UUID

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "block_storage": {
    "hybrid_cache": {
      "disk_count": "6",
      "raid_type": "string",
      "size": "1612709888",
      "used": "26501122"
    },
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "plex0"
      }
    ],
    "primary": {
      "checksum_style": "string",
      "disk_class": "performance",
      "disk_count": "8",
      "disk_type": "string",
      "raid_size": "16",
      "raid_type": "string"
    }
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          }
        }
      }
    ]
  }
}
```

```

        },
        "name": "store1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "used": 0
}
]
},
"create_time": "2018-01-01T12:00:00-04:00",
"dr_home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"home_node": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"inactive_data_reporting": {
    "start_time": "2019-12-12T12:00:00-04:00"
},
"metric": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "duration": "PT15S",
    "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",

```

```

    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"name": "node1_aggr_1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"snaplock_type": "string",
"space": {
  "block_storage": {
    "available": "10156560384",
    "full_threshold_percent": 0,
    "inactive_user_data": "304448",
    "size": "10156769280",
    "used": "2088960"
  },
  "cloud_storage": {
    "used": "402743264"
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "footprint": "608896"
},
"state": "string",
"statistics": {
  "iops_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency_raw": {

```

```

    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "throughput_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "timestamp": "2017-01-25T11:20:13Z"
},
"uuid": "string"
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
787092	The target field cannot be specified for this operation.
7209049	Cannot perform the operation because the aggregate is currently expanding.
8586225	Unexpected error encountered when retrieving metrics and statistics for this aggregate.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
19726390	Unable to provide a recommendation to expand the aggregate.
19726391	Too many unassigned disks visible to the node that owns this aggregate.
19726392	Layout of this aggregate is not a supported configuration.
19726393	Failed to expand the aggregate. Aggregate expansion is not supported on this system.
19726394	Automatic aggregate expansion is not supported on systems with multiple data aggregates.
19726395	Automatic aggregate expansion is not supported when MetroCluster is not configured

Error Code	Description
19726396	Automatic aggregate expansion is not supported when the DR group is not in a normal state
19726397	Aggregates must contain disks with identical disk-types and disk-sizes.
19726538	Cannot perform the operation because the aggregate is not in a healthy state.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
enabled	boolean	Specifies whether the aggregate uses HDDs with SSDs as a cache.
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
_links	_links	
name	string	

primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
checksum_style	string	The checksum style used by the aggregate.
disk_class	string	The class of disks being used by the aggregate.
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
disk_type	string	The type of disk being used by the aggregate.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.

block_storage

Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
hybrid_cache	hybrid_cache	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	mirror	

Name	Type	Description
plexes	array[plex_reference]	Plex reference for each plex in the aggregate.
primary	primary	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud_store

Cloud store

Name	Type	Description
_links	_links	
name	string	
uuid	string	

cloud_storage_tier

Name	Type	Description
cloud_store	cloud_store	Cloud store
used	integer	Capacity used in bytes in the cloud store by this aggregate. This is a cached value calculated every 5 minutes.

cloud_storage

Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.

Name	Type	Description
attach_eligible	boolean	Specifies whether the aggregate is eligible for a cloud store to be attached.
stores	array[cloud_storage_tier]	Configuration information for each cloud storage portion of the aggregate.

Name	Type	Description
tiering_fullness_threshold	integer	The percentage of space in the performance tier that must be used before data is tiered out to the cloud store. Only valid for PATCH operations.

data_encryption

Name	Type	Description
drive_protection_enabled	boolean	Specifies whether the aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	Specifies whether NetApp aggregate encryption is enabled. All data in the aggregate is encrypted.

dr_home_node

Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.

Name	Type	Description
name	string	
uuid	string	

home_node

Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

inactive_data_reporting

Name	Type	Description
enabled	boolean	Specifies whether or not inactive data reporting is enabled on the aggregate.

Name	Type	Description
start_time	string	Timestamp at which inactive data reporting was enabled on the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

The most recent sample of I/O metrics for the aggregate.

Name	Type	Description
_links	_links	
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node

Node where the aggregate currently resides.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

block_storage

Name	Type	Description
available	integer	Space available in bytes.

Name	Type	Description
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either block_storage.inactive_user_data or **.
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for aggregates with a cloud store tier.

efficiency

Storage efficiency

Name	Type	Description
logical_used	integer	Logical used

Name	Type	Description
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

space

Name	Type	Description
block_storage	block_storage	
cloud_storage	cloud_storage	
efficiency	efficiency	Storage efficiency
efficiency_without_snapshots	efficiency_without_snapshots	Storage efficiency that does not include the savings provided by Snapshot copies.
footprint	integer	A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or **.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

The real time I/O statistics for the aggregate.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update an aggregate specified by the UUID

PATCH /storage/aggregates/{uuid}

Introduced In: 9.6

Updates the aggregate specified by the UUID with the properties in the body. This request starts a job and returns a link to that job.

Related ONTAP commands

- `storage aggregate add-disks`
- `storage aggregate mirror`
- `storage aggregate modify`
- `storage aggregate relocation start`
- `storage aggregate rename`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Aggregate UUID
auto_provision_policy	string	query	False	<p>If set to expand, the PATCH operation runs the recommended expansion of the aggregate.</p> <ul style="list-style-type: none"> • Introduced in: 9.8

Name	Type	In	Required	Description
simulate	boolean	query	False	<p>If set to true, the PATCH operation runs a simulated aggregate expansion with the provided input disk count and returns the proposed size of the new aggregate along with any associated warnings.</p> <ul style="list-style-type: none"> • Introduced in: 9.8
disk_size	integer	query	False	<p>If set, PATCH only selects disks of the specified size.</p>

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation

Name	Type	Description
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
name	string	Aggregate name
node	node	Node where the aggregate currently resides.
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

Example request

```
{
  "block_storage": {
    "hybrid_cache": {
      "disk_count": "6",
      "raid_type": "string",
      "size": "1612709888",
      "used": "26501122"
    },
    "mirror": {
      "enabled": "",
      "state": "string"
    },
    "plexes": [
      {
        "name": "plex0"
      }
    ],
    "primary": {
      "disk_count": "8",
      "disk_type": "string",
      "raid_size": "16",
      "raid_type": "string"
    }
  },
  "cloud_storage": {
    "stores": [
      {
        "cloud_store": {
          "name": "store1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        },
        "used": 0
      }
    ]
  },
  "create_time": "2018-01-01T12:00:00-04:00",
  "dr_home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "home_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
}
```

```

"inactive_data_reporting": {
  "start_time": "2019-12-12T12:00:00-04:00"
},
"name": "node1_aggr_1",
"node": {
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "block_storage": {
    "available": "10156560384",
    "full_threshold_percent": 0,
    "inactive_user_data": "304448",
    "size": "10156769280",
    "used": "2088960"
  },
  "cloud_storage": {
    "used": "402743264"
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "footprint": "608896"
},
"state": "string",
"uuid": "string"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	
num_records	integer	Number of records

Name	Type	Description
records	array[aggregate]	
warnings	array[aggregate_warning]	List of validation warnings and remediation advice for the aggregate simulate behavior.

Example response

```
{
  "job": {
    "uuid": "string"
  },
  "records": [
    {
      "block_storage": {
        "hybrid_cache": {
          "disk_count": "6",
          "raid_type": "string",
          "size": "1612709888",
          "used": "26501122"
        },
        "mirror": {
          "enabled": "",
          "state": "string"
        },
        "plexes": [
          {
            "name": "plex0"
          }
        ],
        "primary": {
          "disk_count": "8",
          "disk_type": "string",
          "raid_size": "16",
          "raid_type": "string"
        }
      },
      "cloud_storage": {
        "stores": [
          {
            "cloud_store": {
              "name": "store1",
              "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
            },
            "used": 0
          }
        ]
      },
      "create_time": "2018-01-01T12:00:00-04:00",
      "dr_home_node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
```

```

    },
    "home_node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "inactive_data_reporting": {
      "start_time": "2019-12-12T12:00:00-04:00"
    },
    "name": "node1_aggr_1",
    "node": {
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "space": {
      "block_storage": {
        "available": "10156560384",
        "full_threshold_percent": 0,
        "inactive_user_data": "304448",
        "size": "10156769280",
        "used": "2088960"
      },
      "cloud_storage": {
        "used": "402743264"
      },
      "efficiency": {
        "logical_used": 0,
        "ratio": 0,
        "savings": 0
      },
      "efficiency_without_snapshots": {
        "logical_used": 0,
        "ratio": 0,
        "savings": 0
      },
      "footprint": "608896"
    },
    "state": "string",
    "uuid": "string"
  }
],
"warnings": [
  {
    "action": {
      "arguments": [
        "string"
      ]
    }
  ]
]

```

```

        "message": "string"
    },
    "name": "string",
    "warning": {
        "arguments": [
            "string"
        ],
        "message": "string"
    }
}
]
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
262247	The value is invalid for the field.
460777	Failed to get information on the job.
786434	Cannot connect to node where the aggregate resides.
786435	Internal Error. Failed to create a communication handle.
786439	An aggregate already uses the specified name.
786447	Failed to modify the aggregate.
786456	Failed to add disks to the aggregate.
786458	Failed to rename aggregate.
786468	VLDB is offline.
786472	Node that hosts the aggregate is offline.
786479	Cannot find node ID for the node.
786491	Not enough spares on the node.
786730	Internal Error
786771	Aggregate does not exist.
786787	Aggregate is not online.
786808	Aggregate mirror failed.
786867	Specified aggregate resides on the remote cluster.

Error Code	Description
786911	Not every node in the cluster has the Data ONTAP version required for the feature.
786923	This operation is not allowed during the pre-commit phase of a 7-mode to clustered Data ONTAP transition.
786924	Internal error for an aggregate that is in the pre-commit phase of a 7-mode to clustered Data ONTAP transition.
786955	Modifying raidtype to raid_tec requires a minimum of six disks in the RAID Group.
786956	Modifying raidtype to raid_dp requires a minimum of four disks in the RAID Group.
786965	Spare Selection in userspace failed.
787046	Mirroring of a FabricPool is not allowed.
787092	The target field cannot be specified for this operation.
787144	Aggregate is not a FabricPool.
787156	Modifying the attributes of mirror object store is not allowed.
787169	Only one field can be modified per operation.
787170	Failed to patch the "block_storage.primary.disk_count" because the disk count specified is smaller than existing disk count.
787172	This query is only allowed during the modification of the specified field.
787178	Unmirroring an aggregate with a PATCH operation is not supported.
787187	Internal error. Failed to check if the aggregate is a FabricPool.
1258699	Cannot use all the disks specified for the requested operation.
1263500	Operation will lead to creation of new raid group.
1263501	Operation will exceed half of the maximum volume sizes allowed on the node.
1263502	One spare data partition from at least one of the chosen root-data1-data2 disks will not be used.
1263503	Operation will lead to downsizing of one or more disks.
1263504	Operation will lead to a spares low condition.
1263598	One or more selected disks will be partitioned.

Error Code	Description
1263624	Operation will lead to a no sparecore condition.
2425736	No matching node found for the UUID provided.
7209049	Cannot perform the operation because the aggregate is currently expanding.
7209075	Cannot perform the operation because the volume size limit for this system type would be exceeded.
13108106	Cannot run aggregate relocation because volume expand is in progress.
19726347	There are a number of unassigned disks visible to the node that owns this aggregate.
19726382	Another provisioning operation is in progress on this cluster. Wait a few minutes, and try the operation again.
19726390	Unable to automatically expand this aggregate.
19726391	Too many unassigned disks visible to the node that owns this aggregate.
19726392	Layout of this aggregate is not a supported configuration.
19726393	Failed to expand the aggregate. Aggregate expansion is not supported on this system.
19726394	Automatic aggregate expansion is not supported on systems with multiple data aggregates.
19726395	Automatic aggregate expansion is not supported when MetroCluster is not configured.
19726396	Automatic aggregate expansion is not supported when the DR group is not in a normal state.
19726397	Aggregates must contain disks with identical disk-types and disk-sizes.
19726538	Cannot perform the operation because the aggregate is not in a healthy state.
26542083	Destination node is at higher Data ONTAP version than source node.
26542084	Source node is at higher Data ONTAP version than destination node.
26542097	Unable to get D-blade ID of destination.
26542101	Unable to contact the source node.
26542102	Unable to contact the destination node.
26542120	An SVM migrate operation is in progress. When the migrate operation completes, try the operation again.

Error Code	Description
26542121	A MetroCluster disaster recovery operation is in progress. When the recovery operation completes, try the operation again.
196608334	Failed to modify the aggregate because it contains NAE volumes.
196608335	Failed to modify the aggregate because it contains non-encrypted volumes.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

hybrid_cache

Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.

Name	Type	Description
disk_count	integer	Number of disks used in the cache tier of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
enabled	boolean	Specifies whether the aggregate uses HDDs with SSDs as a cache.
raid_type	string	RAID type for SSD cache of the aggregate. Only provided when hybrid_cache.enabled is 'true'.
size	integer	Total usable space in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.
used	integer	Space used in bytes of SSD cache. Only provided when hybrid_cache.enabled is 'true'.

mirror

Name	Type	Description
enabled	boolean	Aggregate is SyncMirror protected
state	string	

plex_reference

Plex

Name	Type	Description
name	string	

primary

Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

Name	Type	Description
disk_count	integer	Number of disks used in the aggregate. This includes parity disks, but excludes disks in the hybrid cache.
disk_type	string	The type of disk being used by the aggregate.
raid_size	integer	Option to specify the maximum number of disks that can be included in a RAID group.
raid_type	string	RAID type of the aggregate.

block_storage

Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.

Name	Type	Description
hybrid_cache	hybrid_cache	Contains the configuration for the hybrid cache. The hybrid cache is made up of either whole SSDs or storage pool SSDs.
mirror	mirror	
plexes	array[plex_reference]	Plex reference for each plex in the aggregate.
primary	primary	Configuration information for the primary storage portion of the aggregate. This excludes the hybrid cache details.

cloud_store

Cloud store

Name	Type	Description
name	string	
uuid	string	

cloud_storage_tier

Name	Type	Description
cloud_store	cloud_store	Cloud store
used	integer	Capacity used in bytes in the cloud store by this aggregate. This is a cached value calculated every 5 minutes.

cloud_storage

Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.

Name	Type	Description
attach_eligible	boolean	Specifies whether the aggregate is eligible for a cloud store to be attached.
stores	array[cloud_storage_tier]	Configuration information for each cloud storage portion of the aggregate.
tiering_fullness_threshold	integer	The percentage of space in the performance tier that must be used before data is tiered out to the cloud store. Only valid for PATCH operations.

data_encryption

Name	Type	Description
drive_protection_enabled	boolean	Specifies whether the aggregate uses self-encrypting drives with data protection enabled.
software_encryption_enabled	boolean	Specifies whether NetApp aggregate encryption is enabled. All data in the aggregate is encrypted.

dr_home_node

Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.

Name	Type	Description
name	string	
uuid	string	

home_node

Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.

Name	Type	Description
name	string	
uuid	string	

inactive_data_reporting

Name	Type	Description
enabled	boolean	Specifies whether or not inactive data reporting is enabled on the aggregate.
start_time	string	Timestamp at which inactive data reporting was enabled on the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

The most recent sample of I/O metrics for the aggregate.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

node

Node where the aggregate currently resides.

Name	Type	Description
name	string	
uuid	string	

block_storage

Name	Type	Description
available	integer	Space available in bytes.
full_threshold_percent	integer	The aggregate used percentage at which 'monitor.volume.full' EMS is generated.
inactive_user_data	integer	The size that is physically used in the block storage and has a cold temperature, in bytes. This property is only supported if the aggregate is either attached to a cloud store or can be attached to a cloud store. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either <code>block_storage.inactive_user_data</code> or <code>**</code> .
size	integer	Total usable space in bytes, not including WAFL reserve and aggregate Snapshot copy reserve.
used	integer	Space used or reserved in bytes. Includes volume guarantees and aggregate metadata.

cloud_storage

Name	Type	Description
used	integer	Used space in bytes in the cloud store. Only applicable for aggregates with a cloud store tier.

efficiency

Storage efficiency

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

efficiency_without_snapshots

Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

space

Name	Type	Description
block_storage	block_storage	
cloud_storage	cloud_storage	
efficiency	efficiency	Storage efficiency
efficiency_without_snapshots	efficiency_without_snapshots	Storage efficiency that does not include the savings provided by Snapshot copies.

Name	Type	Description
footprint	integer	A summation of volume footprints (including volume guarantees), in bytes. This includes all of the volume footprints in the block_storage tier and the cloud_storage tier. This is an advanced property; there is an added cost to retrieving its value. The field is not populated for either a collection GET or an instance GET unless it is explicitly requested using the <i>fields</i> query parameter containing either footprint or **.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

The real time I/O statistics for the aggregate.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

aggregate

Name	Type	Description
block_storage	block_storage	Configuration information for the locally attached portion of the aggregate. When a cloud store is also used by this aggregate, this is referred to as the performance tier.
cloud_storage	cloud_storage	Configuration information for the cloud storage portion of the aggregate. This is referred to as the capacity tier.
create_time	string	Timestamp of aggregate creation
data_encryption	data_encryption	
dr_home_node	dr_home_node	Node where the aggregate resides after disaster recovery. The value for this field might differ from the 'node' field during switchover.
home_node	home_node	Node where the aggregate resides after giveback. The value for this field might differ from the value of the 'node' field during takeover.
inactive_data_reporting	inactive_data_reporting	
name	string	Aggregate name
node	node	Node where the aggregate currently resides.
space	space	
state	string	Operational state of the aggregate
uuid	string	Aggregate UUID

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

action

Name	Type	Description
arguments	array[string]	Arguments present in the specified action message.
code	integer	Corrective action code of the specified action.
message	string	Specifies the corrective action to be taken to resolve the issue.

warning

Name	Type	Description
arguments	array[string]	Arguments present in the warning message encountered.
code	integer	Warning code of the warning encountered.
message	string	Details of the warning encountered by the aggregate simulate query.

aggregate_warning

Name	Type	Description
action	action	
name	string	Name of the entity that returns the warning.
warning	warning	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Storage aggregate metrics

Storage aggregates UUID metrics endpoint overview

Overview

The Storage Aggregate Metrics API provides historical performance metrics for the specified aggregate. The collection GET operation retrieves read, write, other and total metrics for a given aggregate, in terms of IOPS, latency and throughput. The read and write categories display the I/O operations that service user reads and writes across all the hosted volumes on a given aggregate. The other category encompasses background I/O operations that implement data protection services currently running on the aggregate. IOPs are the number of I/O operations reported per second, throughput is the amount of I/O operations measured in bytes per second and latency is the average response time for an IOP, reported in microseconds. Without a specified time interval, the output is limited to statistics collected at 15 second intervals over the last hour.

Examples

Retrieving metrics for an aggregate

In this example, the API returns a set of records that exist for the aggregate with the given UUID for the last hour.


```
# The API:
/api/storage/aggregates/{uuid}/metrics

#The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-
bad0-005056b48388/metrics?max_records=4" -H "accept: application/json"

#The response:
{
  "records": [
    {
      "timestamp": "2019-01-14T23:33:45Z"
    },
    {
      "timestamp": "2019-01-14T23:33:30Z"
    },
    {
      "timestamp": "2019-01-14T23:33:15Z"
    },
    {
      "timestamp": "2019-01-14T23:33:00Z"
    }
  ],
  "num_records": 4
}
```

Retrieving metrics for an aggregate with a set timestamp

In this example, the API returns metric values for latency, IOPS, and throughput properties such as read, write and total. The status and duration for which the metrics are requested are also returned.

```

#The API:
/api/storage/aggregates/{uuid}/metrics?timestamp={timestamp}

#The call:
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-bad0-005056b48388/metrics?timestamp=2019-01-1T23:33:00Z" -H "accept: application/json"

#The response:
{
  "records": [
    {
      "uuid": "538bf337-1b2c-11e8-bad0-005056b48388",
      "timestamp": "2019-01-01T23:33:00Z",
      "status": "ok",
      "duration": "PT15S",
      "throughput": {
        "read": 6826,
        "write": 205892,
        "other": 0,
        "total": 212718
      },
      "latency": {
        "read": 148,
        "write": 216,
        "other": 0,
        "total": 199
      },
      "iops": {
        "read": 1,
        "write": 5,
        "other": 0,
        "total": 6
      }
    }
  ]
}

```

Retrieving metrics for an aggregate for a set time interval

In this example, the API returns the requested metrics for the given time interval of 1 week. The interval value can be

1 hour, 1 day, 1 week, 1 month or 1 year. If the interval value is not set, a default value of 1 hour is used.

```
#The API:
```

```
/api/storage/aggregates/{uuid}/metrics
```

#The call:

```
curl -X GET "https://<mgmt-ip>/api/storage/aggregates/538bf337-1b2c-11e8-bad0-005056b48388/metrics?return_timeout=15&fields=*&interval=1w&max_records=4"
-H "accept: application/json"
```

#The response:

```
{
  "records": [
    {
      "timestamp": "2019-01-01T23:30:00Z",
      "status": "ok",
      "duration": "PT30M",
      "throughput": {
        "read": 268328,
        "write": 5556255,
        "other": 0,
        "total": 5824584
      },
      "latency": {
        "read": 156,
        "write": 430,
        "other": 0,
        "total": 318
      },
      "iops": {
        "read": 18,
        "write": 26,
        "other": 0,
        "total": 45
      }
    },
    {
      "timestamp": "2019-01-01T23:00:00Z",
      "status": "ok",
      "duration": "PT30M",
      "throughput": {
        "read": 474266,
        "write": 6121908,
        "other": 0,
        "total": 6596175
      },
      "latency": {
        "read": 154,
```

```

        "write": 448,
        "other": 0,
        "total": 262
    },
    "iops": {
        "read": 48,
        "write": 28,
        "other": 0,
        "total": 76
    }
},
{
    "timestamp": "2019-01-01T22:30:00Z",
    "status": "ok",
    "duration": "PT30M",
    "throughput": {
        "read": 540164,
        "write": 2411356,
        "other": 26244685,
        "total": 29196206
    },
    "latency": {
        "read": 159,
        "write": 394,
        "other": 192,
        "total": 193
    },
    "iops": {
        "read": 94,
        "write": 16,
        "other": 437,
        "total": 548
    }
},
{
    "timestamp": "2019-01-01T22:00:00Z",
    "status": "ok",
    "duration": "PT30M",
    "throughput": {
        "read": 2842,
        "write": 2765407,
        "other": 0,
        "total": 2768249
    },
    "latency": {
        "read": 189,

```

```

    "write": 540,
    "other": 0,
    "total": 523
  },
  "iops": {
    "read": 0,
    "write": 13,
    "other": 0,
    "total": 13
  }
},
"num_records": 4
}

```

Related ONTAP commands

- `statistics aggregate show`

Retrieve historical performance metrics for an aggregate

GET /storage/aggregates/{uuid}/metrics

Introduced In: 9.7

Retrieves historical performance metrics for an aggregate.

Parameters

Name	Type	In	Required	Description
timestamp	string	query	False	Filter by timestamp
duration	string	query	False	Filter by duration
throughput.total	integer	query	False	Filter by throughput.total
throughput.read	integer	query	False	Filter by throughput.read
throughput.other	integer	query	False	Filter by throughput.other
throughput.write	integer	query	False	Filter by throughput.write
latency.total	integer	query	False	Filter by latency.total

Name	Type	In	Required	Description
latency.read	integer	query	False	Filter by latency.read
latency.other	integer	query	False	Filter by latency.other
latency.write	integer	query	False	Filter by latency.write
status	string	query	False	Filter by status
iops.total	integer	query	False	Filter by iops.total
iops.read	integer	query	False	Filter by iops.read
iops.other	integer	query	False	Filter by iops.other
iops.write	integer	query	False	Filter by iops.write
uuid	string	path	True	Unique identifier of the aggregate.

Name	Type	In	Required	Description
interval	string	query	False	<p>The time range for the data. Values can be 1h, 1d, 1w, 1m, or 1y. The period for each time range is as follows:</p> <ul style="list-style-type: none"> • 1h: Metrics over the most recent hour sampled over 15 seconds. • 1d: Metrics over the most recent day sampled over 5 minutes. • 1w: Metrics over the most recent week sampled over 30 minutes. • 1m: Metrics over the most recent month sampled over 2 hours. • 1y: Metrics over the most recent year sampled over a day. • Default value: 1 • enum: ["1h", "1d", "1w", "1m", "1y"]

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8586225	Encountered unexpected error in retrieving metrics for the requested aggregate.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

records

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Report cluster-wide storage details across different tiers

GET `/storage/cluster`

Introduced In: 9.6

Reports cluster wide storage details across different tiers. By default, this endpoint returns all fields. Supports the following roles: admin, and readonly.

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
block_storage	block_storage	
cloud_storage	cloud_storage	
efficiency	space_efficiency	
efficiency_without_snapshots	space_efficiency	

Example response

```
{
  "block_storage": {
    "medias": [
      {
        "efficiency": {
          "logical_used": 0,
          "ratio": 0,
          "savings": 0
        },
        "type": "string"
      }
    ]
  },
  "cloud_storage": {
    "used": 0
  },
  "efficiency": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  },
  "efficiency_without_snapshots": {
    "logical_used": 0,
    "ratio": 0,
    "savings": 0
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

space_efficiency

Name	Type	Description
logical_used	integer	Logical used
ratio	number	Data reduction ratio (logical_used / used)
savings	integer	Space saved by storage efficiencies (logical_used - used)

medias

Name	Type	Description
available	integer	Available space
efficiency	space_efficiency	
size	integer	Total space
type	string	The type of media being used
used	integer	Used space

block_storage

Name	Type	Description
inactive_data	integer	Inactive data across all aggregates
medias	array[medias]	
size	integer	Total space across the cluster
used	integer	Space used (includes volume reserves)

cloud_storage

Name	Type	Description
used	integer	Total space used in cloud.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage disks

Storage disks endpoint overview

Retrieving storage disk information

The storage disk GET API retrieves all of the disks in the cluster.

Examples

1) Retrieve a list of disks from the cluster.

The following example shows the response with a list of disks in the cluster:

```
# The API:
/api/storage/disks

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/disks" -H "accept:
application/hal+json"

# The response:
{
  "records": [
```

```

{
  "name": "1.24.4",
  "_links": {
    "self": {
      "href": "/api/storage/disks/1.24.4"
    }
  }
},
{
  "name": "1.24.3",
  "_links": {
    "self": {
      "href": "/api/storage/disks/1.24.3"
    }
  }
},
{
  "name": "1.24.5",
  "_links": {
    "self": {
      "href": "/api/storage/disks/1.24.5"
    }
  }
},
{
  "name": "1.24.0",
  "_links": {
    "self": {
      "href": "/api/storage/disks/1.24.0"
    }
  }
},
{
  "name": "1.24.2",
  "_links": {
    "self": {
      "href": "/api/storage/disks/1.24.2"
    }
  }
},
{
  "name": "1.24.1",
  "_links": {
    "self": {
      "href": "/api/storage/disks/1.24.1"
    }
  }
}

```

```

    }
  }
],
"num_records": 6,
"_links": {
  "self": {
    "href": "/api/storage/disks"
  }
}
}
}

```

2) Retrieve a specific disk from the cluster.

The following example shows the response of the requested disk. If there is no disk with the requested name, an error is returned.

```

# The API:
/api/storage/disks/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/disks/1.24.3" -H "accept:
application/hal+json"

# The response:
{
  "name": "1.24.3",
  "uid":
    "50000394:0808AA88:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000",
  "serial_number": "EC47PC5021SW",
  "model": "X421_FAL12450A10",
  "vendor": "NETAPP",
  "firmware_version": "NA02",
  "usable_size": 438304768000,
  "rpm": 10000,
  "type": "sas",
  "class": "performance",
  "container_type": "aggregate",
  "pool": "pool0",
  "state": "present",
  "node": {
    "uuid": "3a89ed49-8c6d-11e8-93bc-00a0985a64b6",
    "name": "node-2",

```

```

    "_links": {
      "self": {
        "href": "/api/cluster/nodes/3a89ed49-8c6d-11e8-93bc-00a0985a64b6"
      }
    },
    "home_node": {
      "uuid": "3a89ed49-8c6d-11e8-93bc-00a0985a64b6",
      "name": "node-2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/3a89ed49-8c6d-11e8-93bc-00a0985a64b6"
        }
      }
    },
    "aggregates": [
      {
        "uuid": "3fd9c345-ba91-4949-a7b1-6e2b898d74e3",
        "name": "node_2_SAS_1",
        "_links": {
          "self": {
            "href": "/api/storage/aggregates/3fd9c345-ba91-4949-a7b1-6e2b898d74e3"
          }
        }
      }
    ],
    "shelf": {
      "uid": "10318311901725526608",
      "_links": {
        "self": {
          "href": "/api/storage/shelves/10318311901725526608"
        }
      }
    },
    "bay": 3,
    "_links": {
      "self": {
        "href": "/api/storage/disks/1.24.3"
      }
    }
  }
}

```

Modifying storage disk

The storage disk PATCH API modifies disk ownership or encrypting drive authentication keys (AKs) in the cluster.

Examples

Rekey the data AK of all encrypting drives to an AK selected automatically by the system.

```
# The API:
/api/storage/disks

# The call:
curl -X PATCH "https://<mgmt-
ip>/api/storage/disks?name=&encryption_operation=rekey_data_auto_id" -H
"accept: application/hal+json" -H "Content-Type: application/hal+json"

# The response contains the number of disks attempted.
{
  "num_records": 32
}
```

Retrieve a collection of disks

GET /storage/disks

Introduced In: 9.6

Retrieves a collection of disks.

Related ONTAP commands

- storage disk show

Learn more

- [DOC /storage/disks](#)

Parameters

Name	Type	In	Required	Description
dr_node.uuid	string	query	False	Filter by dr_node.uuid

Name	Type	In	Required	Description
dr_node.name	string	query	False	Filter by dr_node.name
aggregates.uuid	string	query	False	Filter by aggregates.uuid
aggregates.name	string	query	False	Filter by aggregates.name
container_type	string	query	False	Filter by container_type
protection_mode	string	query	False	Filter by protection_mode <ul style="list-style-type: none"> Introduced in: 9.7
state	string	query	False	Filter by state
serial_number	string	query	False	Filter by serial_number
bay	integer	query	False	Filter by bay
shelf.uuid	string	query	False	Filter by shelf.uuid
usable_size	integer	query	False	Filter by usable_size
firmware_version	string	query	False	Filter by firmware_version
uid	string	query	False	Filter by uid
class	string	query	False	Filter by class
type	string	query	False	Filter by type
name	string	query	False	Filter by name
rpm	integer	query	False	Filter by rpm
model	string	query	False	Filter by model

Name	Type	In	Required	Description
key_id.fips	string	query	False	Filter by key_id.fips • Introduced in: 9.7
key_id.data	string	query	False	Filter by key_id.data • Introduced in: 9.7
rated_life_used_percent	integer	query	False	Filter by rated_life_used_percent
home_node.name	string	query	False	Filter by home_node.name
home_node.uuid	string	query	False	Filter by home_node.uuid
drawer.id	integer	query	False	Filter by drawer.id
drawer.slot	integer	query	False	Filter by drawer.slot
fips_certified	boolean	query	False	Filter by fips_certified • Introduced in: 9.7
self_encrypting	boolean	query	False	Filter by self_encrypting • Introduced in: 9.7
pool	string	query	False	Filter by pool
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
vendor	string	query	False	Filter by vendor
fields	array[string]	query	False	Specify the fields to return.

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[disk]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "bay": "1",
      "class": "solid_state",
      "container_type": "spare",
      "dr_node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "encryption_operation": "string",
      "firmware_version": "NA51",
      "home_node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "key_id": {
        "data": "string",
        "fips": "string"
      }
    },
  ],
}
```

```

"model": "X421_HCOBE450A10",
"name": "1.0.1",
"node": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"pool": "pool0",
"protection_mode": "string",
"rated_life_used_percent": "10",
"rpm": "15000",
"serial_number": "KHG2VX8R",
"shelf": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "uid": "7777841915827391056"
},
"state": "present",
"type": "ssd",
"uid":
"002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000",
"usable_size": "959934889984",
"vendor": "NETAPP"
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

dr_node

Name	Type	Description
name	string	
uuid	string	

drawer

Name	Type	Description
id	integer	
slot	integer	

home_node

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	
uuid	string	

key_id

Name	Type	Description
data	string	Key ID of the data authentication key
fips	string	Key ID of the FIPS authentication key

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

shelf

Name	Type	Description
_links	_links	
uid	string	

disk

Name	Type	Description
aggregates	array[aggregates]	List of aggregates sharing this disk
bay	integer	Disk shelf bay
class	string	Disk class
container_type	string	Type of overlying disk container
dr_node	dr_node	
drawer	drawer	

Name	Type	Description
encryption_operation	string	Encryption operation to apply to the drives. Possible values are: <ul style="list-style-type: none"> • rekey_data_default • rekey_data_auto_id
fips_certified	boolean	
firmware_version	string	
home_node	home_node	
key_id	key_id	
model	string	
name	string	Cluster-wide disk name
node	node	
pool	string	Pool to which disk is assigned
protection_mode	string	Mode of drive data protection and FIPS compliance. Possible values are: <ul style="list-style-type: none"> • <i>open</i> - Disk is unprotected • <i>data</i> - Data protection only without FIPS compliance • <i>part</i> - Partial protection with FIPS compliance only • <i>full</i> - Full data and FIPS compliance protection
rated_life_used_percent	integer	Percentage of rated life used
rpm	integer	Revolutions per minute
self_encrypting	boolean	
serial_number	string	
shelf	shelf	
state	string	State
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	

Name	Type	Description
vendor	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update disk ownership, change authentication keys, or sanitize disks

PATCH /storage/disks

Introduced In: 9.7

Updates disk ownership or authentication keys.

Related ONTAP commands

- `storage disk assign`
- `storage encryption disk modify -data-key-id`
- `security key-manager key query -key-type NSE-AK`

Learn more

- [DOC /storage/disks](#)

Parameters

Name	Type	In	Required	Description
name	string	query	False	Disk name

Name	Type	In	Required	Description
encryption_operation	string	query	False	Name of the operation to apply to encrypting disks. <ul style="list-style-type: none"> enum: ["rekey_data_default", "rekey_data_auto_id"]
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none"> Default value:

Request Body

Name	Type	Description
aggregates	array[aggregates]	List of aggregates sharing this disk
bay	integer	Disk shelf bay
class	string	Disk class
container_type	string	Type of overlying disk container
drawer	drawer	
encryption_operation	string	Encryption operation to apply to the drives. Possible values are: <ul style="list-style-type: none"> rekey_data_default rekey_data_auto_id
fips_certified	boolean	
firmware_version	string	
key_id	key_id	
model	string	
name	string	Cluster-wide disk name
pool	string	Pool to which disk is assigned

Name	Type	Description
protection_mode	string	Mode of drive data protection and FIPS compliance. Possible values are: <ul style="list-style-type: none"> • <i>open</i> - Disk is unprotected • <i>data</i> - Data protection only without FIPS compliance • <i>part</i> - Partial protection with FIPS compliance only • <i>full</i> - Full data and FIPS compliance protection
rated_life_used_percent	integer	Percentage of rated life used
rpm	integer	Revolutions per minute
self_encrypting	boolean	
serial_number	string	
state	string	State
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	
vendor	string	

Example request

```
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "bay": "1",
  "class": "solid_state",
  "container_type": "spare",
  "encryption_operation": "string",
  "firmware_version": "NA51",
  "key_id": {
    "data": "string",
    "fips": "string"
  },
  "model": "X421_HCOBE450A10",
  "name": "1.0.1",
  "pool": "pool0",
  "protection_mode": "string",
  "rated_life_used_percent": "10",
  "rpm": "15000",
  "serial_number": "KHG2VX8R",
  "state": "present",
  "type": "ssd",
  "uid":
    "002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000",
  "usable_size": "959934889984",
  "vendor": "NETAPP"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
721066	Node is outside the list of controllers for disk.
1441795	Setting the data key ID to the manufacture secure ID is not allowed when in FIPS-compliance mode.
14155777	The operation failed on one or more disks.
14155778	No self-encrypting disks were specified.
14155779	Status from a node shows that a conflicting operation has occurred. Some disk controls might have changed.
14155780	Could not retrieve the required key ID from the key manager.
14155786	Changes to encryption controls are not allowed with drive assignment.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

aggregates

Aggregate

Name	Type	Description
name	string	
uuid	string	

dr_node

Name	Type	Description
name	string	
uuid	string	

drawer

Name	Type	Description
id	integer	
slot	integer	

home_node

Name	Type	Description
name	string	
uuid	string	

key_id

Name	Type	Description
data	string	Key ID of the data authentication key
fips	string	Key ID of the FIPS authentication key

node

Name	Type	Description
name	string	
uuid	string	

shelf

Name	Type	Description
uid	string	

disk

Name	Type	Description
aggregates	array[aggregates]	List of aggregates sharing this disk
bay	integer	Disk shelf bay
class	string	Disk class
container_type	string	Type of overlying disk container
drawer	drawer	
encryption_operation	string	Encryption operation to apply to the drives. Possible values are: <ul style="list-style-type: none"> rekey_data_default rekey_data_auto_id
fips_certified	boolean	
firmware_version	string	
key_id	key_id	
model	string	
name	string	Cluster-wide disk name
pool	string	Pool to which disk is assigned

Name	Type	Description
protection_mode	string	Mode of drive data protection and FIPS compliance. Possible values are: <ul style="list-style-type: none"> • <i>open</i> - Disk is unprotected • <i>data</i> - Data protection only without FIPS compliance • <i>part</i> - Partial protection with FIPS compliance only • <i>full</i> - Full data and FIPS compliance protection
rated_life_used_percent	integer	Percentage of rated life used
rpm	integer	Revolutions per minute
self_encrypting	boolean	
serial_number	string	
state	string	State
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	
vendor	string	

Retrieve a specific disk

GET /storage/disks/{name}

Introduced In: 9.6

Retrieves a specific disk.

Related ONTAP commands

- `storage disk show`

Learn more

- [DOC /storage/disks](#)

Parameters

Name	Type	In	Required	Description
name	string	path	True	Disk name
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
aggregates	array[aggregates]	List of aggregates sharing this disk
bay	integer	Disk shelf bay
class	string	Disk class
container_type	string	Type of overlying disk container
dr_node	dr_node	
drawer	drawer	
encryption_operation	string	Encryption operation to apply to the drives. Possible values are: <ul style="list-style-type: none">• rekey_data_default• rekey_data_auto_id
fips_certified	boolean	
firmware_version	string	
home_node	home_node	
key_id	key_id	
model	string	
name	string	Cluster-wide disk name
node	node	
pool	string	Pool to which disk is assigned

Name	Type	Description
protection_mode	string	Mode of drive data protection and FIPS compliance. Possible values are: <ul style="list-style-type: none"> • <i>open</i> - Disk is unprotected • <i>data</i> - Data protection only without FIPS compliance • <i>part</i> - Partial protection with FIPS compliance only • <i>full</i> - Full data and FIPS compliance protection
rated_life_used_percent	integer	Percentage of rated life used
rpm	integer	Revolutions per minute
self_encrypting	boolean	
serial_number	string	
shelf	shelf	
state	string	State
type	string	Disk interface type
uid	string	The unique identifier for a disk
usable_size	integer	
vendor	string	

Example response

```
{
  "aggregates": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "bay": "1",
  "class": "solid_state",
  "container_type": "spare",
  "dr_node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "encryption_operation": "string",
  "firmware_version": "NA51",
  "home_node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "key_id": {
    "data": "string",
    "fips": "string"
  },
  "model": "X421_HCOBE450A10",
  "name": "1.0.1",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

```
},
"pool": "pool0",
"protection_mode": "string",
"rated_life_used_percent": "10",
"rpm": "15000",
"serial_number": "KHG2VX8R",
"shelf": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "uid": "7777841915827391056"
},
"state": "present",
"type": "ssd",
"uid":
"002538E5:71B00B2F:00000000:00000000:00000000:00000000:00000000:00000000:00000000:00000000",
"usable_size": "959934889984",
"vendor": "NETAPP"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

dr_node

Name	Type	Description
name	string	
uuid	string	

drawer

Name	Type	Description
id	integer	
slot	integer	

home_node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

key_id

Name	Type	Description
data	string	Key ID of the data authentication key
fips	string	Key ID of the FIPS authentication key

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

shelf

Name	Type	Description
_links	_links	
uid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage storage file clones

Create a clone of the file

POST /storage/file/clone

Introduced In: 9.6

Creates a clone of the file.

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
autodelete	boolean	Mark clone file for auto deletion.
destination_path	string	Relative path of the clone/destination file in the volume.
is_backup	boolean	Mark clone file for backup.
overwrite_destination	boolean	Destination file gets overwritten.
range	array[string]	List of block ranges for sub-file cloning in the format "source-file-block-number:destination-file-block-number:block-count"
source_path	string	Relative path of the source file in the volume.
volume	volume	

Example request

```
{
  "destination_path": "dest_file1, dir1/dest_file2",
  "range": [
    "10:10:5",
    "20:20:10"
  ],
  "source_path": "src_file1, dir1/src_file2,
  ../.snapshot/snap1/src_file3",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

file_clone

File clone

Name	Type	Description
autodelete	boolean	Mark clone file for auto deletion.
destination_path	string	Relative path of the clone/destination file in the volume.
is_backup	boolean	Mark clone file for backup.
overwrite_destination	boolean	Destination file gets overwritten.
range	array[string]	List of block ranges for sub-file cloning in the format "source-file-block-number:destination-file-block-number:block-count"
source_path	string	Relative path of the source file in the volume.

Name	Type	Description
volume	volume	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Start a file copy operation

POST /storage/file/copy

Introduced In: 9.8

Starts a file copy operation.

Required properties

- `files_to_copy` - List of files with the destination they are to be copied to.

Default property values

- `cutover_time` - 10

- `hold_quiescence` - *false*
- `max_throughput` - *0*
- `reference_cutover_time` - *10*

Related ONTAP commands

- `volume file copy start`

Examples

Copying two files

The POST request is used to copy file(s).

```
# The API:
/api/storage/file/copy

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/file/copy" -H "accept:
application/hal+json" -d
'{"files_to_copy":[{"source":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},
"path":"d1/src_f1"},"destination":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},
"path":"d1/dst_f1"}},
{"source":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},"path":"d1/src_f
2"},"destination":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},"path":"
d1/dst_f2"}}}]}'

# The response:
{
  "job": {
    "uuid": "b89bc5dd-94a3-11e8-a7a3-0050568edf84",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/b89bc5dd-94a3-11e8-a7a3-0050568edf84"
      }
    }
  }
}
```

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
cutover_time	integer	The maximum amount of time (in seconds) that the source can be quiesced before a destination file must be made available for read-write traffic.

Name	Type	Description
files_to_copy	array[files_to_copy]	A list of source files along with the destinations they are copied to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory.
hold_quiescence	boolean	Specifies whether the source file should be held quiescent for the duration of the copy operation.
max_throughput	integer	The maximum amount of data (in bytes) that can be transferred per second in support of this operation.
reference_cutover_time	integer	The maximum amount of time (in seconds) that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic.
reference_file	reference_file	

Example request

```
{
  "cutover_time": "10",
  "files_to_copy": [
    {
      "destination": {
        "path": "string",
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      },
      "source": {
        "path": "string",
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    },
    {
      "reference_cutover_time": "10",
      "reference_file": {
        "path": "string",
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    }
  ]
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
7012352	File locations are inconsistent. All files must be on the same volume.
7012353	Exceeded the file operations supported number of files.
7012354	Unable to pair the number of source files to destination files.
7012357	Cannot start a file operation until all cluster nodes support the file operations capability.
7012358	The specified source path is invalid.
7012359	The specified destination path is invalid.
7012360	The SVMs are not in an intracluster peering relationship.
7012361	The SVMs peering relationship does not include application "file-copy".
7012362	The SVMs are not yet in a peered state yet.
7012363	Cannot copy files. All file operations must be managed by the destination SVM's administrator.
7012365	Copying a file between clusters is not supported.
7012367	A reference path may only be specified if multiple source paths are specified.
7012368	The reference path must have a matching source path.

Error Code	Description
7012371	The reference cutover time exceeds the maximum allowable time.
7012374	Source volume and destination volume have different home clusters.
7012376	Operation not allowed on a volume that is part of a SnapMirror Synchronous relationship.
7012377	Cannot start a file copy operation on the volume because an active volume conversion is in progress.
196608143	Cannot start operation. The volume is undergoing a secure purge operation.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

file_reference

Name	Type	Description
path	string	Path of the file or directory.
svm	svm	
volume	volume	

files_to_copy

Name	Type	Description
destination	file_reference	
source	file_reference	

reference_file

Name	Type	Description
path	string	The source reference file. If a reference file is specified, data for other files being copied will be transferred as a difference from the reference file. This can save bandwidth and destination storage if the specified source files share blocks. If provided, this input must match one of the source file paths. This input need not be provided if only one source file is specified.
volume	volume	

file_copy

File copy

Name	Type	Description
cutover_time	integer	The maximum amount of time (in seconds) that the source can be quiesced before a destination file must be made available for read-write traffic.
files_to_copy	array[files_to_copy]	A list of source files along with the destinations they are copied to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory.
hold_quiescence	boolean	Specifies whether the source file should be held quiescent for the duration of the copy operation.
max_throughput	integer	The maximum amount of data (in bytes) that can be transferred per second in support of this operation.

Name	Type	Description
reference_cutover_time	integer	The maximum amount of time (in seconds) that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic.
reference_file	reference_file	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Start a file move operation

POST /storage/file/move

Introduced In: 9.8

Starts a file move operation.

Required properties

- `files_to_move` - List of files with the destination they are to be moved to.

Default property values

- `cutover_time` - *10*
- `hold_quiescence` - *false*
- `max_throughput` - *0*
- `reference_cutover_time` - *10*

Related ONTAP commands

- `volume file move start`

Examples

Copying two files

The POST request is used to move file(s).

```
# The API:
/api/storage/file/move

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/file/move" -H "accept:
application/hal+json" -d
'{"files_to_move":[{"source":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},
"path":"d1/src_f1"},"destination":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},
"path":"d1/dst_f1"}},
{"source":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},"path":"d1/src_f
2"},"destination":{"volume":{"name":"vol_a"},"svm":{"name":"vs0"},"path":"
d1/dst_f2"}}}]}'

# The response:
{
  "job": {
    "uuid": "b89bc5dd-94a3-11e8-a7a3-0050568edf84",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/b89bc5dd-94a3-11e8-a7a3-0050568edf84"
      }
    }
  }
}
```

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
cutover_time	integer	The maximum amount of time (in seconds) that the source can be quiesced before a destination file must be made available for read-write traffic.
files_to_move	array[files_to_move]	A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory.
max_throughput	integer	The maximum amount of data (in bytes) that can be transferred per second in support of this operation.
reference_cutover_time	integer	The maximum amount of time (in seconds) that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic.
reference_file	reference_file	

Example request

```
{
  "cutover_time": "10",
  "files_to_move": [
    {
      "destination": {
        "path": "string",
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      },
      "source": {
        "path": "string",
        "svm": {
          "name": "svm1",
          "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
        },
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    },
    {
      "reference_cutover_time": "10",
      "reference_file": {
        "path": "string",
        "volume": {
          "name": "volume1",
          "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
        }
      }
    }
  ]
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
7012352	File locations are inconsistent. All files must be on the same volume.
7012353	Exceeded the file operations supported number of files.
7012354	Unable to pair the number of source files to destination files.
7012357	Cannot start a file operation until all cluster nodes support the file operations capability.
7012358	The specified source path is invalid.
7012359	The specified destination path is invalid.
7012360	The SVMs are not in an intracenter peering relationship.
7012361	The SVMs peering relationship does not include application "file-move".
7012362	The SVMs are not yet in a peered state yet.
7012363	Cannot move files. All file operations must be managed by the destination SVM's administrator.
7012365	Copying a file between clusters is not supported.
7012367	A reference path may only be specified if multiple source paths are specified.
7012368	The reference path must have a matching source path.

Error Code	Description
7012371	The reference cutover time exceeds the maximum allowable time.
7012374	Source volume and destination volume have different home clusters.
7012376	Operation not allowed on a volume that is part of a SnapMirror Synchronous relationship.
7012377	Cannot start a file move operation on the volume because an active volume conversion is in progress.
196608143	Cannot start the operation. The volume is undergoing a secure purge operation.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7

destination

Name	Type	Description
path	string	Path of the file or directory.
svm	svm	SVM, applies only to SVM-scoped objects.
volume	volume	

source

Name	Type	Description
path	string	Path of the file or directory.
svm	svm	SVM, applies only to SVM-scoped objects.
volume	volume	

files_to_move

Name	Type	Description
destination	destination	
source	source	

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

reference_file

Name	Type	Description
path	string	The source reference file. If a reference file is specified, data for other files being moved will be transferred as a difference from the reference file. This can save bandwidth and destination storage if the specified source files share blocks. If provided, this input must match one of the source file paths. This input need not be provided if only one source file is specified.
volume	volume	

file_move

File move

Name	Type	Description
cutover_time	integer	The maximum amount of time (in seconds) that the source can be quiesced before a destination file must be made available for read-write traffic.
files_to_move	array[files_to_move]	A list of source files along with the destination file they are moved to. If the terminal path component of the destination is a directory, then the source file's basename is replicated in that directory.
max_throughput	integer	The maximum amount of data (in bytes) that can be transferred per second in support of this operation.
reference_cutover_time	integer	The maximum amount of time (in seconds) that the source reference file can be quiesced before the corresponding destination file must be made available for read-write traffic.
reference_file	reference_file	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage FlexCache volumes

Storage flexcache flexcaches endpoint overview

Overview

FlexCache is a persistent cache of an origin volume. An origin volume can only be a FlexVol while a FlexCache is always a FlexGroup.

The following relationship configurations are supported:

- – Intra-Vserver where FlexCache and the corresponding origin volume reside in the same Vserver.

- – Cross-Vserver but intra-cluster where FlexCache and the origin volume reside in the same cluster but belong to different Vservers.

- – Cross-cluster where FlexCache and the origin volume reside in different clusters.

FlexCache supports fan-out and more than one FlexCache can be created from one origin volume. This API retrieves and manages FlexCache configurations in the cache cluster.

FlexCache APIs

The following APIs can be used to perform operations related with FlexCache:

- – GET /api/storage/flexcache/flexcaches

- – GET /api/storage/flexcache/flexcaches/{uuid}

- – POST /api/storage/flexcache/flexcaches

- – DELETE /api/storage/flexcache/flexcaches/{uuid}

Examples

Creating a FlexCache

The POST request is used to create a FlexCache.

```
# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/flexcache/flexcaches" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
  \"aggregates\": [ { \"name\": \"aggr_1\" } ], \"name\": \"fc_333\",
  \"origins\": [ { \"svm\": { \"name\": \"vs_3\" }, \"volume\": {
    \"name\": \"vol_o1\" } } ], \"svm\": { \"name\": \"vs_1\" } }"

# The response:
{
  "job": {
    "uuid": "e751dd5d-0f3c-11e9-8b2b-0050568e0b79",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/e751dd5d-0f3c-11e9-8b2b-0050568e0b79"
      }
    }
  }
}
```

Retrieving FlexCache attributes

The GET request is used to retrieve FlexCache attributes. The object includes a large set of fields which can be expensive to retrieve. Most notably, the fields `size`, `guarantee.type`, `aggregates`, `path`, `origins.ip_address`, `origins.size`, and `origins.state` are expensive to retrieve. The recommended method to use this API is to filter and retrieve only the required fields.

```
# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/flexcaches?" -H
"accept: application/json"

# The response:
{
  "records": [
    {
      "uuid": "04d5e07b-0ebe-11e9-8180-0050568e0b79",
      "name": "fc_322",
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/flexcaches/04d5e07b-0ebe-11e9-
```



```

8180-0050568e0b79"
    }
  }
},
{
  "uuid": "47902654-0ea4-11e9-8180-0050568e0b79",
  "name": "fc_321",
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/flexcaches/47902654-0ea4-11e9-
8180-0050568e0b79"
    }
  }
},
{
  "uuid": "77e911ff-0ebe-11e9-8180-0050568e0b79",
  "name": "fc_323",
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/flexcaches/77e911ff-0ebe-11e9-
8180-0050568e0b79"
    }
  }
},
{
  "uuid": "ddb42bbc-0e95-11e9-8180-0050568e0b79",
  "name": "fc_32",
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/flexcaches/ddb42bbc-0e95-11e9-
8180-0050568e0b79"
    }
  }
},
{
  "uuid": "ec774932-0f3c-11e9-8b2b-0050568e0b79",
  "name": "fc_333",
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-
8b2b-0050568e0b79"
    }
  }
}
],
"num_records": 5,

```

```

    "_links": {
      "self": {
        "href": "/api/storage/flexcache/flexcaches?"
      }
    }
  }
}

```

Retrieving the attributes of a FlexCache

The GET request is used to retrieve the attributes of a FlexCache. The object includes a large set of fields which can be expensive to retrieve. Most notably, the fields `size`, `guarantee.type`, `aggregates`, `path`, `origins.ip_address`, `origins.size`, and `origins.state` are expensive to retrieve. The recommended method to use this API is to filter and retrieve only the required fields.

```

# The API:
/api/storage/flexcache/flexcaches/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79" -H "accept: application/json"

# The response:
{
  "uuid": "ec774932-0f3c-11e9-8b2b-0050568e0b79",
  "name": "fc_333",
  "svm": {
    "name": "vs_1",
    "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
  },
  "size": 4294967296,
  "guarantee": {
    "type": "volume"
  },
  "aggregates": [
    {
      "name": "aggr_1",
      "uuid": "26f34b76-88f8-4a47-b5e0-d8e901fb1114"
    }
  ],
  "origins": [
    {
      "ip_address": "10.140.103.175",
      "size": 20971520,
      "create_time": "2019-01-03T15:19:55+05:30",
      "state": "online",
      "volume": {

```

```

        "name": "vol_o1",
        "uuid": "2bc957dd-2617-4afb-8d2f-66ac6070d313"
    },
    "svm": {
        "name": "vs_3",
        "uuid": "8aa2cd28-0e92-11e9-b391-0050568e4115"
    },
    "cluster": {
        "name": "node2",
        "uuid": "50733f81-0e90-11e9-b391-0050568e4115"
    }
}
],
"_links": {
    "self": {
        "href": "/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79"
    }
}
}

```

Deleting a FlexCache

The DELETE request is used to delete a FlexCache.

```

# The API:
/api/storage/flexcache/flexcaches

# The call:
curl -X DELETE "https://<mgmt-
ip>/api/storage/flexcache/flexcaches/ec774932-0f3c-11e9-8b2b-0050568e0b79"
-H "accept: application/json"

# The response:
{
    "job": {
        "uuid": "e17994f2-0f3e-11e9-8b2b-0050568e0b79",
        "_links": {
            "self": {
                "href": "/api/cluster/jobs/e17994f2-0f3e-11e9-8b2b-0050568e0b79"
            }
        }
    }
}

```

Retrieve a FlexCache volume in the cluster

GET /storage/flexcache/flexcaches

Introduced In: 9.6

Retrieves FlexCache in the cluster.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `origins.ip_address` - IP address of origin.
- `origins.size` - Physical size of origin.
- `origins.state` - State of origin.
- `size` - Physical size of FlexCache.
- `guarantee.type` - Space guarantee style of FlexCache.
- `aggregates.name` or `aggregates.uuid` - Name or UUID of aggregate of FlexCache volume.
- `path` - Fully-qualified path of the owning SVM's namespace where the FlexCache is mounted.

Related ONTAP commands

- `volume flexcache show`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
aggregates.uuid	string	query	False	Filter by aggregates.uuid
aggregates.name	string	query	False	Filter by aggregates.name
constituents_per_aggregate	integer	query	False	Filter by constituents_per_aggregate
origins.svm.uuid	string	query	False	Filter by origins.svm.uuid

Name	Type	In	Required	Description
origins.svm.name	string	query	False	Filter by origins.svm.name
origins.state	string	query	False	Filter by origins.state
origins.ip_address	string	query	False	Filter by origins.ip_address
origins.cluster.name	string	query	False	Filter by origins.cluster.name
origins.cluster.uuid	string	query	False	Filter by origins.cluster.uuid
origins.size	integer	query	False	Filter by origins.size
origins.volume.uuid	string	query	False	Filter by origins.volume.uuid
origins.volume.name	string	query	False	Filter by origins.volume.name
origins.create_time	string	query	False	Filter by origins.create_time
name	string	query	False	Filter by name
lock_consistent_when_disconnected	boolean	query	False	Filter by lock_consistent_when_disconnected <ul style="list-style-type: none"> Introduced in: 9.8
size	integer	query	False	Filter by size
path	string	query	False	Filter by path
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name

Name	Type	In	Required	Description
use_tiered_aggregate	boolean	query	False	Filter by use_tiered_aggregate <ul style="list-style-type: none"> Introduced in: 9.8
guarantee.type	string	query	False	Filter by guarantee.type <ul style="list-style-type: none"> Introduced in: 9.7
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none"> Default value: 1 Max value: 120 Min value: 0
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[flexcache]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "guarantee": {
        "type": "string"
      },
      "name": "vol1",
      "origins": [
        {
          "cluster": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "create_time": "2018-06-04T19:00:00Z",
          "ip_address": "10.10.10.7",

```



```

    "size": 0,
    "state": "string",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ],
  "path": "/user/my_fc",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

guarantee

Name	Type	Description
type	string	The type of space guarantee of this volume in the aggregate.

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

svm

FlexCache SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

flexcache

Defines the cache endpoint of FlexCache.

Name	Type	Description
_links	_links	
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
guarantee	guarantee	
lock_consistent_when_disconnected	boolean	Specifies if the origin will honor the cache side locks when doing the lock checks in the disconnected mode. The default value will be set to false to allow origin to skip such lock checks with the disconnected cache and go ahead with processing of the conflicting locks. If set to true the origin denies the lock operation (open or byte-range) if it conflicts with the authority granted to the disconnected cache.
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.

Name	Type	Description
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The use_tiered_aggregate is only used when auto-provisioning a FlexCache volume.
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a FlexCache volume in the cluster

POST /storage/flexcache/flexcaches

Introduced In: 9.6

Creates a FlexCache in the cluster.

Required properties

- `name` - Name of FlexCache volume.
- `origins.volume.name` or `origins.volume.uuid` - Name or UUID of origin volume.
- `origins.svm.name` - Name of origin Vserver.
- `svm.name` or `svm.uuid` - Name or UUID of Vserver where FlexCache will be created.

Recommended optional properties

- `path` - Path to mount the FlexCache volume

Default property values

If not specified in POST, the following default property values are assigned:

- `size` - 10% of origin volume size or 1GB per constituent, whichever is greater.
- `guarantee.type` - Same as for a non-FlexCache FlexGroup volume.
- `constituents_per_aggregate` - 4 if `aggregates.name` or `aggregates.uuid` is used.
- `use_tiered_aggregate` - false if `aggr-list` is not used. This property is only used when auto-provisioning a FlexCache volume.
- `lock_consistent_when_disconnected` - false. This property specifies if the origin will honor the cache side locks when doing the lock checks in the disconnected mode.

Related ONTAP commands

- `volume flexcache create`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
aggregates	array[aggregates]	
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.

Name	Type	Description
guarantee	guarantee	
lock_consistent_when_disconnected	boolean	Specifies if the origin will honor the cache side locks when doing the lock checks in the disconnected mode. The default value will be set to false to allow origin to skip such lock checks with the disconnected cache and go ahead with processing of the conflicting locks. If set to true the origin denies the lock operation (open or byte-range) if it conflicts with the authority granted to the disconnected cache.
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The use_tiered_aggregate is only used when auto-provisioning a FlexCache volume.
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

Example request

```
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "guarantee": {
    "type": "string"
  },
  "name": "vol1",
  "origins": [
    {
      "cluster": {
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04T19:00:00Z",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "string",
      "svm": {
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ],
  "path": "/user/my_fc",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66846870	Either the SVM name or origin volume name is missing
66846871	Constituents per aggregate are specified but aggregate name is missing
66846872	More than one origin volume is specified
66846873	The specified SVM UUID is incorrect for the specified SVM name
66846874	The specified aggregate UUID is incorrect for the specified aggregate name
66846875	The specified aggregate name does not exist
66846876	The specified SVM does not exist or is not peered
66846877	The specified origin SVM name is of zero length
66846878	The specified SVM UUID is invalid
66846730	Failed to create a FlexCache volume
66846760	The specified SVM is not a data Vserver
66846787	The specified aggregate is a SnapLock aggregate
66846812	The specified aggregate is a Composite aggregate

Error Code	Description
66846812	The specified junction path is under a FlexCache volume
66846834	FlexCache encryption requires a cluster version of 9.6 or higher
66846835	A volume encryption license is not found

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

aggregates

Aggregate

Name	Type	Description
name	string	
uuid	string	

guarantee

Name	Type	Description
type	string	The type of space guarantee of this volume in the aggregate.

cluster

Name	Type	Description
name	string	
uuid	string	

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

svm

FlexCache SVM

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

flexcache

Defines the cache endpoint of FlexCache.

Name	Type	Description
aggregates	array[aggregates]	

Name	Type	Description
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
guarantee	guarantee	
lock_consistent_when_disconnected	boolean	Specifies if the origin will honor the cache side locks when doing the lock checks in the disconnected mode. The default value will be set to false to allow origin to skip such lock checks with the disconnected cache and go ahead with processing of the conflicting locks. If set to true the origin denies the lock operation (open or byte-range) if it conflicts with the authority granted to the disconnected cache.
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The use_tiered_aggregate is only used when auto-provisioning a FlexCache volume.

Name	Type	Description
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a FlexCache volume

```
DELETE /storage/flexcache/flexcaches/{uuid}
```

Introduced In: 9.6

Deletes a FlexCache. If a FlexCache volume is online, it is offlined before deletion.

Related ONTAP commands

- `volume flexcache delete`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of the FlexCache.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
66846879	The specified volume UUID is not a FlexCache volume
66846731	Failed to delete the FlexCache volume
524546	Failed to delete the FlexCache volume because the FlexCache volume is not unmounted

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve attributes of the FlexCache volume in the cluster

GET /storage/flexcache/flexcaches/{uuid}

Introduced In: 9.6

Retrieves attributes of the FlexCache in the cluster.

Expensive properties

There is an added cost to retrieving values for these properties. They are included by default in GET. The recommended method to use this API is to filter and retrieve only the required fields. See [Requesting specific fields](#) to learn more.

- `origins.ip_address` - IP address of origin.
- `origins.size` - Physical size of origin.
- `origins.state` - State of origin.
- `size` - Physical size of FlexCache.
- `guarantee.type` - Space guarantee style of FlexCache.
- `aggregates.name` or `aggregates.uuid` - Name or UUID of aggregate of FlexCache volume.
- `path` - Fully-qualified path of the owning SVM's namespace where the FlexCache is mounted.

Related ONTAP commands

- `volume flexcache show`

Learn more

- [DOC /storage/flexcache/flexcaches](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of FlexCache.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
<code>_links</code>	_links	
<code>aggregates</code>	array[aggregates]	

Name	Type	Description
constituents_per_aggregate	integer	Number of FlexCache constituents per aggregate when the 'aggregates' field is mentioned.
guarantee	guarantee	
lock_consistent_when_disconnected	boolean	Specifies if the origin will honor the cache side locks when doing the lock checks in the disconnected mode. The default value will be set to false to allow origin to skip such lock checks with the disconnected cache and go ahead with processing of the conflicting locks. If set to true the origin denies the lock operation (open or byte-range) if it conflicts with the authority granted to the disconnected cache.
name	string	FlexCache name
origins	array[flexcache_relationship]	
path	string	The fully-qualified path in the owning SVM's namespace at which the FlexCache is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one FlexCache be mounted at any given junction path.
size	integer	Physical size of the FlexCache. The recommended size for a FlexCache is 10% of the origin volume. The minimum FlexCache constituent size is 1GB.
svm	svm	FlexCache SVM
use_tiered_aggregate	boolean	Specifies whether or not a Fabricpool-enabled aggregate can be used in FlexCache creation. The use_tiered_aggregate is only used when auto-provisioning a FlexCache volume.

Name	Type	Description
uuid	string	FlexCache UUID. Unique identifier for the FlexCache.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "guarantee": {
    "type": "string"
  },
  "name": "vol1",
  "origins": [
    {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04T19:00:00Z",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
```



```

    },
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ],
  "path": "/user/my_fc",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

guarantee

Name	Type	Description
type	string	The type of space guarantee of this volume in the aggregate.

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

svm

FlexCache SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage storage FlexCache origin volumes

Storage flexcache origins endpoint overview

Overview

FlexCache is a persistent cache of an origin volume. An origin volume can only be a FlexVol while a FlexCache is always a FlexGroup.

The following relationship configurations are supported:

- – Intra-Vserver where FlexCache and the corresponding origin volume reside in the same Vserver.
- – Cross-Vserver but intra-cluster where FlexCache and the origin volume reside in the same cluster but belong to different Vservers.
- – Cross-cluster where FlexCache and the origin volume reside in different clusters.

FlexCache supports fan-out and more than one FlexCache can be created from one origin volume. This API retrieves the origin of FlexCache onfigurations in the origin cluster.

FlexCache APIs

The following APIs can be used to perform operations related to the origin of a FlexCache:

- – GET /api/storage/flexcache/origins
- – GET /api/storage/flexcache/origins/{uuid}

Examples

Retrieving origins of FlexCache attributes

The GET request is used to retrieve the origins of FlexCache attributes.

```
# The API:
/api/storage/flexcache/origins

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/origins?" -H
"accept: application/json"

# The response:
{
  "records": [
    {
      "uuid": "2bc957dd-2617-4afb-8d2f-66ac6070d313",
      "name": "vol_01",
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/origins/2bc957dd-2617-4afb-8d2f-66ac6070d313"
        }
      }
    },
    {
      "uuid": "80fcaee4-0dc2-488b-afb8-86d28a34cda8",
      "name": "vol_1",
      "_links": {
        "self": {
          "href": "/api/storage/flexcache/origins/80fcaee4-0dc2-488b-afb8-86d28a34cda8"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/origins?"
    }
  }
}
```

Retrieving the attributes of an origin volume

The GET request is used to retrieve the attributes of an origin volume.

```
# The API:
/api/storage/flexcache/origins/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/flexcache/origins/80fcaee4-0dc2-488b-afb8-86d28a34cda8" -H "accept: application/json"

# The response:
{
  "uuid": "80fcaee4-0dc2-488b-afb8-86d28a34cda8",
  "name": "vol_1",
  "svm": {
    "name": "vs_3",
    "uuid": "8aa2cd28-0e92-11e9-b391-0050568e4115"
  },
  "flexcaches": [
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-02T19:27:22+05:30",
      "volume": {
        "name": "fc_42",
        "uuid": "4e7f9d49-0e96-11e9-aed0-0050568eddbbe"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddbbe"
      },
      "cluster": {
        "name": "node4",
        "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbbe"
      }
    },
    {
      "ip_address": "10.140.103.183",
      "create_time": "2019-01-02T21:08:34+05:30",
      "volume": {
        "name": "fc_421",
        "uuid": "71ee8f36-0ea4-11e9-aed0-0050568eddbbe"
      },
      "svm": {
        "name": "vs_1_4",
        "uuid": "36f68322-0e93-11e9-aed0-0050568eddbbe"
      }
    }
  ]
}
```

```

    },
    "cluster": {
      "name": "node4",
      "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbe"
    }
  },
  {
    "ip_address": "10.140.103.183",
    "create_time": "2019-01-03T11:14:38+05:30",
    "volume": {
      "name": "fc_422"
    },
    "svm": {
      "name": "vs_1_4",
      "uuid": "36f68322-0e93-11e9-aed0-0050568eddbe"
    },
    "cluster": {
      "name": "node4",
      "uuid": "c32f16b8-0e90-11e9-aed0-0050568eddbe"
    }
  },
  {
    "ip_address": "10.140.103.179",
    "size": 4294967296,
    "create_time": "2019-01-02T19:24:14+05:30",
    "state": "online",
    "volume": {
      "name": "fc_32",
      "uuid": "ddb42bbc-0e95-11e9-8180-0050568e0b79"
    },
    "svm": {
      "name": "vs_1",
      "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
    },
    "cluster": {
      "name": "node3",
      "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
    }
  },
  {
    "ip_address": "10.140.103.179",
    "size": 4294967296,
    "create_time": "2019-01-02T21:07:23+05:30",
    "state": "online",
    "volume": {
      "name": "fc_321",

```



```

    "uuid": "47902654-0ea4-11e9-8180-0050568e0b79"
  },
  "svm": {
    "name": "vs_1",
    "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
  },
  "cluster": {
    "name": "node3",
    "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
  }
},
{
  "ip_address": "10.140.103.179",
  "size": 4294967296,
  "create_time": "2019-01-03T00:11:38+05:30",
  "state": "online",
  "volume": {
    "name": "fc_322",
    "uuid": "04d5e07b-0ebe-11e9-8180-0050568e0b79"
  },
  "svm": {
    "name": "vs_1",
    "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
  },
  "cluster": {
    "name": "node3",
    "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
  }
},
{
  "ip_address": "10.140.103.179",
  "size": 4294967296,
  "create_time": "2019-01-03T00:14:52+05:30",
  "state": "online",
  "volume": {
    "name": "fc_323",
    "uuid": "77e911ff-0ebe-11e9-8180-0050568e0b79"
  },
  "svm": {
    "name": "vs_1",
    "uuid": "e708fbe2-0e92-11e9-8180-0050568e0b79"
  },
  "cluster": {
    "name": "node3",
    "uuid": "8eb21b3b-0e90-11e9-8180-0050568e0b79"
  }
}

```

```

    }
  ],
  "_links": {
    "self": {
      "href": "/api/storage/flexcache/origins/80fcaee4-0dc2-488b-afb8-86d28a34cda8"
    }
  }
}

```

Retrieve the origin of a FlexCache volume in the cluster

GET /storage/flexcache/origins

Introduced In: 9.6

Retrieves origin of FlexCache in the cluster.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `flexcaches.ip_address` - IP address of FlexCache.
- `flexcaches.size` - Physical size of FlexCache.
- `flexcaches.guarantee.type` - Space guarantee style of FlexCache.
- `flexcaches.state` - State of FlexCache.

Related ONTAP commands

- `volume flexcache origin show-caches`

Learn more

- [DOC /storage/flexcache/origins](#)

Parameters

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name

Name	Type	In	Required	Description
flexcaches.svm.uuid	string	query	False	Filter by flexcaches.svm.uuid
flexcaches.svm.name	string	query	False	Filter by flexcaches.svm.name
flexcaches.state	string	query	False	Filter by flexcaches.state
flexcaches.ip_address	string	query	False	Filter by flexcaches.ip_address
flexcaches.cluster.name	string	query	False	Filter by flexcaches.cluster.name
flexcaches.cluster.uuid	string	query	False	Filter by flexcaches.cluster.uuid
flexcaches.size	integer	query	False	Filter by flexcaches.size
flexcaches.volume.uuid	string	query	False	Filter by flexcaches.volume.uuid
flexcaches.volume.name	string	query	False	Filter by flexcaches.volume.name
flexcaches.create_time	string	query	False	Filter by flexcaches.create_time
name	string	query	False	Filter by name

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[flexcache_origin]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "flexcaches": [
        {
          "cluster": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "cluster1",
            "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
          },
          "create_time": "2018-06-04T19:00:00Z",
          "ip_address": "10.10.10.7",
          "size": 0,
          "state": "string",
          "svm": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            },
            "name": "svm1",
            "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
          },
          "volume": {
            "_links": {
              "self": {
                "href": "/api/resourcelink"
              }
            }
          }
        }
      ]
    }
  ]
}
```

```

        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  ],
  "name": "vol1, vol_2",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

svm

Origin volume SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

flexcache_origin

Defines the origin endpoint of FlexCache.

Name	Type	Description
_links	_links	
flexcaches	array[flexcache_relationship]	

Name	Type	Description
name	string	Origin volume name
svm	svm	Origin volume SVM
uuid	string	Origin volume UUID. Unique identifier for origin of FlexCache.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve attributes of the origin of a FlexCache volume in the cluster

GET /storage/flexcache/origins/{uuid}

Introduced In: 9.6

Retrieves attributes of the origin of a FlexCache in the cluster.

Expensive properties

There is an added cost to retrieving values for these properties. They are included by default in GET results. The recommended method to use this API is to filter and retrieve only the required fields. See [Requesting specific fields](#) to learn more.

- `flexcaches.ip_address` - IP address of FlexCache.
- `flexcaches.size` - Physical size of FlexCache.
- `flexcaches.guarantee.type` - Space guarantee style of FlexCache.

- `flexcaches.state` - State of FlexCache.

Related ONTAP commands

- `volume flexcache origin show-caches`

Learn more

- [DOC /storage/flexcache/origins](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of origin of FlexCache.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
flexcaches	array[flexcache_relationship]	
name	string	Origin volume name
svm	svm	Origin volume SVM
uuid	string	Origin volume UUID. Unique identifier for origin of FlexCache.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "flexcaches": [
    {
      "cluster": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "cluster1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "create_time": "2018-06-04T19:00:00Z",
      "ip_address": "10.10.10.7",
      "size": 0,
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    },
    {
      "name": "vol1, vol_2",
      "svm": {
        "_links": {
```

```
    "self": {
      "href": "/api/resourcelink"
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563512"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

cluster

Name	Type	Description
_links	_links	
name	string	
uuid	string	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

flexcache_relationship

Name	Type	Description
cluster	cluster	
create_time	string	Creation time of the relationship.
ip_address	string	Cluster management IP of the remote cluster.
size	integer	Size of the remote volume.
state	string	Volume state
svm	svm	
volume	volume	

svm

Origin volume SVM

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

View and create monitored files

Storage monitored-files endpoint overview

Retrieving all monitored files

```
# The API:
GET /api/storage/monitored-files

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/monitored-files"
```

Provisioning a monitored file

```
# The API:
POST /api/storage/monitored-files

# The call:
curl -d "@test_mfiles_post.txt" -X POST "https://<mgmt-ip>/api/storage/monitored-files"
test_mfiles_post.txt (body):
{
  "svm": {
    "name": "vs0"
  },
  "volume": {
    "name": "vol1"
  },
  "path": "/a/b/c/file.txt"
}
```

Removing a file from the monitored files list

```
# The API:
DELETE /api/storage/monitored-files/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/monitored-files/6f68c85b-45e1-11e9-8fc7-005056bbc848"
```

Alternate method for removing files from the monitored files list

Monitored files can also be deleted via a combination of any of (uuid, svm.name, svm.uuid, volume.name, volume.uuid, path).

For example, to remove all monitored-files from monitoring in a single svm named vs0, use the following

```
# The API:
DELETE /api/storage/monitored-files

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/monitored-
files?svm.name=vs0"
```

Performance monitoring

Performance of the monitored file can be monitored by the `metric.*` and `statistics.*` properties. These fields show the performance of the monitored file in terms of IOPS, latency and throughput. The `metric.*` properties denote an average whereas `statistics.*` properties denote a real-time monotonically increasing value aggregated across all nodes.

Retrieve all monitored files

GET /storage/monitored-files

Introduced In: 9.8

Retrieves all monitored files.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `statistics.*`
- `metric.*`

Learn more

- [DOC /storage/monitored-files](#)

Parameters

Name	Type	In	Required	Description
statistics.iops_raw.to tal	integer	query	False	Filter by statistics.iops_raw.to tal
statistics.iops_raw.re ad	integer	query	False	Filter by statistics.iops_raw.r ead

Name	Type	In	Required	Description
statistics.iops_raw.other	integer	query	False	Filter by statistics.iops_raw.other
statistics.iops_raw.write	integer	query	False	Filter by statistics.iops_raw.write
statistics.latency_raw.total	integer	query	False	Filter by statistics.latency_raw.total
statistics.latency_raw.read	integer	query	False	Filter by statistics.latency_raw.read
statistics.latency_raw.other	integer	query	False	Filter by statistics.latency_raw.other
statistics.latency_raw.write	integer	query	False	Filter by statistics.latency_raw.write
statistics.timestamp	string	query	False	Filter by statistics.timestamp
statistics.status	string	query	False	Filter by statistics.status
statistics.throughput_raw.total	integer	query	False	Filter by statistics.throughput_raw.total
statistics.throughput_raw.read	integer	query	False	Filter by statistics.throughput_raw.read
statistics.throughput_raw.other	integer	query	False	Filter by statistics.throughput_raw.other
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write
path	string	query	False	Filter by path

Name	Type	In	Required	Description
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
metric.timestamp	string	query	False	Filter by metric.timestamp
metric.duration	string	query	False	Filter by metric.duration
metric.throughput.total	integer	query	False	Filter by metric.throughput.total
metric.throughput.read	integer	query	False	Filter by metric.throughput.read
metric.throughput.other	integer	query	False	Filter by metric.throughput.other
metric.throughput.write	integer	query	False	Filter by metric.throughput.write
metric.latency.total	integer	query	False	Filter by metric.latency.total
metric.latency.read	integer	query	False	Filter by metric.latency.read
metric.latency.other	integer	query	False	Filter by metric.latency.other
metric.latency.write	integer	query	False	Filter by metric.latency.write
metric.status	string	query	False	Filter by metric.status
metric.iops.total	integer	query	False	Filter by metric.iops.total

Name	Type	In	Required	Description
metric.iops.read	integer	query	False	Filter by metric.iops.read
metric.iops.other	integer	query	False	Filter by metric.iops.other
metric.iops.write	integer	query	False	Filter by metric.iops.write
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
uuid	string	query	False	Filter by uuid
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of files provisioned for monitoring.
records	array[monitored_file]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "metric": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    },
    {
      "path": "/a/b/c/file.txt",
      "statistics": {
        "iops_raw": {
```

```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
    "name": "svm1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
},
"uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
"volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

performance_metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

performance_metric_raw_reference

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Name	Type	Description
name	string	The name of the svm.
uuid	string	The UUID of the svm.

volume

Name	Type	Description
name	string	The name of the volume.

Name	Type	Description
uuid	string	The UUID of the volume.

monitored_file

Name	Type	Description
_links	_links	
metric	performance_metric	Performance numbers, such as IOPS latency and throughput.
path	string	Path of the file to be monitored.
statistics	performance_metric_raw_reference	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
svm	svm	
uuid	string	Unique identifier created for identifying the file that is monitored.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Create a monitored file

POST /storage/monitored-files

Introduced In: 9.8

Creates a monitored file.

Required properties

- `svm.name` - SVM where the file to be monitored exists.
- `volume.name` - Volume where the file to be monitored exists.
- `path` - Path to the file to be monitored.

Learn more

- [DOC /storage/monitored-files](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
path	string	Path of the file to be monitored.
svm	svm	
uuid	string	Unique identifier created for identifying the file that is monitored.
volume	volume	

Example request

```
{
  "path": "/a/b/c/file.txt",
  "svm": {
    "name": "svm1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of files provisioned for monitoring.
records	array[monitored_file]	

Example response

```
{
  "records": [
    {
      "path": "/a/b/c/file.txt",
      "svm": {
        "name": "svm1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "volume": {
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

performance_metric

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.

Name	Type	Description
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Performance metric for write I/O operations.

performance_metric_raw_reference

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Name	Type	Description
name	string	The name of the svm.
uuid	string	The UUID of the svm.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	The UUID of the volume.

monitored_file

Name	Type	Description
path	string	Path of the file to be monitored.
svm	svm	
uuid	string	Unique identifier created for identifying the file that is monitored.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a specific file

DELETE /storage/monitored-files/{uuid}

Introduced In: 9.8

Removes the file from the list of monitored files.

Learn more

- [DOC /storage/monitored-files](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	

Response

Status: 200, Ok

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieves historical performance metrics for the monitored file

GET /storage/monitored-files/{monitored_file.uuid}/metrics

Introduced In: 9.8

Retrieves historical performance metrics for the monitored file.

Parameters

Name	Type	In	Required	Description
throughput.total	integer	query	False	Filter by throughput.total
throughput.read	integer	query	False	Filter by throughput.read
throughput.other	integer	query	False	Filter by throughput.other
throughput.write	integer	query	False	Filter by throughput.write
timestamp	string	query	False	Filter by timestamp
duration	string	query	False	Filter by duration
status	string	query	False	Filter by status
iops.total	integer	query	False	Filter by iops.total
iops.read	integer	query	False	Filter by iops.read
iops.other	integer	query	False	Filter by iops.other
iops.write	integer	query	False	Filter by iops.write
latency.total	integer	query	False	Filter by latency.total
latency.read	integer	query	False	Filter by latency.read
latency.other	integer	query	False	Filter by latency.other
latency.write	integer	query	False	Filter by latency.write
monitored_file.uuid	string	path	True	Unique identifier of the monitored file.

Name	Type	In	Required	Description
interval	string	query	False	<p>The time range for the data. Examples can be 1w, 1m, 1y. The period for each time range is as follows:</p> <ul style="list-style-type: none"> • 1w: Metrics over the most recent week sampled over 30 minutes. • 1m: Metrics over the most recent month sampled over 2 hours. • 1y: Metrics over the most recent year sampled over a day. • Default value: 1 • enum: ["1w", "1m", "1y"]
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

records

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage storage ports

Storage ports endpoint overview

Retrieving storage port information

The storage port GET API retrieves all of the storage ports in the cluster.

Examples

1) Retrieve a list of storage ports from the cluster

The following example shows the response with a list of storage ports in the cluster:

```
# The API:
/api/storage/ports

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/ports" -H "accept:
application/hal+json"

# The response:
{
  "records": [
```

```

{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
      }
    }
  },
  "name": "0a",
  "_links": {
    "self": {
      "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0a"
    }
  }
},
{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
      }
    }
  },
  "name": "0b",
  "_links": {
    "self": {
      "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0b"
    }
  }
},
{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
      }
    }
  }
}

```

```

    }
  },
  "name": "0c",
  "_links": {
    "self": {
      "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0c"
    }
  }
},
{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
      }
    }
  },
  "name": "0d",
  "_links": {
    "self": {
      "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0d"
    }
  }
},
{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
      }
    }
  },
  "name": "0e",
  "_links": {
    "self": {
      "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0e"
    }
  }
}

```

```

    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
        }
      }
    },
    "name": "0f",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0f"
      }
    }
  },
  {
    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
        }
      }
    },
    "name": "0g",
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0g"
      }
    }
  },
],
"num_records": 7,
"_links": {
  "self": {
    "href": "/api/storage/ports"
  }
}

```

```
}  
}  
}
```

2) Retrieve a specific storage port from the cluster

The following example shows the response of the requested storage port. If there is no storage port with the requested node uuid and name, an error is returned.

```
# The API:
/api/storage/ports/{node.uuid}/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0a" -H "accept: application/hal+json"

# The response:
{
  "node": {
    "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
    "name": "node-1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
      }
    }
  },
  "name": "0a",
  "description": "SAS Host Adapter 0a (PMC-Sierra PM8001 rev. C)",
  "wwn": "500a098003633df0",
  "speed": 6,
  "cable": {
    "part_number": "112-00429+A0",
    "serial_number": "629230774",
    "identifier": "500a0980066e2c01-500a098003633df0",
    "length": "0.5m"
  },
  "state": "online",
  "_links": {
    "self": {
      "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0a"
    }
  }
}
```

Retrieve storage ports

GET /storage/ports

Introduced In: 9.6

Retrieves a collection of storage ports.

Related ONTAP commands

- `storage port show`

Learn more

- [DOC /storage/ports](#)

Parameters

Name	Type	In	Required	Description
board_name	string	query	False	Filter by board_name
mac_address	string	query	False	Filter by mac_address
description	string	query	False	Filter by description
name	string	query	False	Filter by name
error.message	string	query	False	Filter by error.message
error.corrective_action	string	query	False	Filter by error.corrective_action
speed	number	query	False	Filter by speed
cable.serial_number	string	query	False	Filter by cable.serial_number
cable.identifier	string	query	False	Filter by cable.identifier
cable.part_number	string	query	False	Filter by cable.part_number
cable.length	string	query	False	Filter by cable.length
wwn	string	query	False	Filter by wwn
state	string	query	False	Filter by state
serial_number	string	query	False	Filter by serial_number

Name	Type	In	Required	Description
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
mode	string	query	False	Filter by mode <ul style="list-style-type: none"> Introduced in: 9.8
part_number	string	query	False	Filter by part_number
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none"> Default value: 1 Max value: 120 Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[storage_port]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "board_name": "string",
      "cable": {
        "identifier": "500a0980000b6c3f-50000d1703544b80",
        "length": "2m",
        "part_number": "112-00431+A0",
        "serial_number": "616930439"
      },
      "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
      "error": {
        "corrective_action": "string",
        "message": "string"
      },
      "mac_address": "string",
      "mode": "storage",
      "name": "2a",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "part_number": "111-03801",
      "serial_number": "7A2463CC45B",
      "speed": "6",
      "state": "online",
      "wnn": "50000d1703544b80"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

error

Name	Type	Description
corrective_action	string	Error corrective action
message	string	Error message

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

storage_port

Name	Type	Description
board_name	string	
cable	cable	
description	string	
error	error	
mac_address	string	
mode	string	Operational mode of a non-dedicated Ethernet port
name	string	
node	node	
part_number	string	
serial_number	string	
speed	number	Operational port speed in Gbps
state	string	
wwn	string	World Wide Name

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a storage port

GET /storage/ports/{node.uuid}/{name}

Introduced In: 9.6

Retrieves a specific storage port.

Related ONTAP commands

- `storage port show`

Learn more

- [DOC /storage/ports](#)

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	Node UUID
name	string	path	True	Port name
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
board_name	string	
cable	cable	
description	string	
error	error	
mac_address	string	
mode	string	Operational mode of a non-dedicated Ethernet port
name	string	
node	node	
part_number	string	
serial_number	string	

Name	Type	Description
speed	number	Operational port speed in Gbps
state	string	
wwn	string	World Wide Name

Example response

```
{
  "board_name": "string",
  "cable": {
    "identifier": "500a0980000b6c3f-50000d1703544b80",
    "length": "2m",
    "part_number": "112-00431+A0",
    "serial_number": "616930439"
  },
  "description": "SAS Host Adapter 2a (PMC-Sierra PM8072 rev. C)",
  "error": {
    "corrective_action": "string",
    "message": "string"
  },
  "mac_address": "string",
  "mode": "storage",
  "name": "2a",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  },
  "name": "node1",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
  "part_number": "111-03801",
  "serial_number": "7A2463CC45B",
  "speed": "6",
  "state": "online",
  "wwn": "50000d1703544b80"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

error

Name	Type	Description
corrective_action	string	Error corrective action
message	string	Error message

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage QoS policies

Storage Qos policies endpoint overview

Quality of Service Configuration

A QoS policy defines measurable service level objectives (SLOs) that apply to the storage objects with which the policy is associated. There are two types of policies that can be configured: fixed, which defines a fixed SLO, or adaptive which defines a variable SLO for a storage object. Adaptive policies vary the SLO depending on the space usage of the storage object. A policy can be either a fixed policy or an adaptive one, not both.

Service level objectives include minimum and maximum limits on throughput in terms of IOPS. Only maximum limits can be set in terms of both IOPS and/or throughput (MB/s). A QoS policy can be used to enforce SLOs for multiple storage objects by specifying "capacity_shared" to true. For example, if a QoS policy with "capacity_shared" is set to true and it has maximum_throughput_iops set to 1000, and this policy is assigned to four volumes, then the combined throughput of all four volumes is limited to 1000 IOPS. If "capacity_shared" is set to false then, each storage object will have it's SLOs enforced individually. For example, in the previous case if the same policy was applied to four volumes but with "capacity_shared" set to false, then each of the volumes would be limited to 1000 IOPS individually. Once "capacity_shared" is set, it cannot be modified.

Adaptive parameters can specify the variable SLOs in terms of IOPS/TB. The actual IOPS enforced on the storage object can be calculated using the allocated space on the storage object. The policies are enforced individually amongst storage objects.

Examples

1) Create a fixed QoS policy

The following example shows how to create a fixed QoS policy to limit throughput for a storage object between 5000 IOPS and 10000 IOPS which has capacity_shared set to false. This QoS policy can be used as a template to apply on multiple storage objects to provide individual SLOs to each object.

+

```
curl -X POST
"https://172.21.69.245/api/storage/qos/policies?return_timeout=0" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"fixed\": { \"capacity_shared\": false, \"max_throughput_iops\": 10000,
\"min_throughput_iops\": 5000 }, \"name\":
\"qos_policy_5000_to_10000_iops\", \"svm\": { \"name\": \"vs0\" } }"
```

2) Create an adaptive QoS policy

The following example shows how to create an adaptive QoS policy which provides 5000 IOPS per GB of allocated space for a storage object with a peak of 6000 IOPS. Minimum IOPS regardless of allocated space are 1000 IOPS.

+

```
curl -X POST
"https://172.21.69.245/api/storage/qos/policies?return_timeout=0" -H
"accept: application/json" -H "Content-Type: application/json" -d "{
\"adaptive\": { \"absolute_min_iops\": 1000, \"expected_iops\": 5000,
\"peak_iops\": 6000 }, \"name\": \"adaptive_pg_5k_to_6k\", \"svm\": {
\"name\": \"vs0\" } }"
```

-

3) Update an existing QoS policy

The following example shows how to update SLOs of an existing QoS policy and also rename it.

+

```
curl -X PATCH "https://172.21.69.245/api/storage/qos/policies/d38bafc0-
5a51-11e9-bd5b-005056ac6f1f?return_timeout=0" -H "accept:
application/json" -H "Content-Type: application/json" -d "{ \"fixed\": {
\"max_throughput_iops\": 15000, \"min_throughput_iops\": 10000 },
\"name\": \"qos_policy_10k_to_15k_iops\" }"
```

4) Delete an existing QoS policy

When a QoS policy is deleted any associations of the policy with a storage objects are also removed.

+

```
curl -X DELETE "https://172.21.69.245/api/storage/qos/policies/d38bafc0-5a51-11e9-bd5b-005056ac6f1f?return_timeout=0" -H "accept: application/json"
```

Retrieve QoS policies

GET /storage/qos/policies

Introduced In: 9.6

Retrieves a collection of QoS policies.

Parameters

Name	Type	In	Required	Description
fixed.min_throughput_mbps	integer	query	False	Filter by fixed.min_throughput_mbps <ul style="list-style-type: none">Introduced in: 9.8
fixed.min_throughput_iops	integer	query	False	Filter by fixed.min_throughput_iops
fixed.max_throughput_iops	integer	query	False	Filter by fixed.max_throughput_iops
fixed.capacity_shared	boolean	query	False	Filter by fixed.capacity_shared
fixed.max_throughput_mbps	integer	query	False	Filter by fixed.max_throughput_mbps
adaptive.peak_iops	integer	query	False	Filter by adaptive.peak_iops

Name	Type	In	Required	Description
adaptive.absolute_min_iops	integer	query	False	Filter by adaptive.absolute_min_iops
adaptive.expected_iops	integer	query	False	Filter by adaptive.expected_iops
object_count	integer	query	False	Filter by object_count
name	string	query	False	Filter by name
uuid	string	query	False	Filter by uuid
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
error	error	
num_records	integer	Number of records
records	array[qos_policy]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "extreme",
      "object_count": 0,
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

_links

Name	Type	Description
self	href	

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

Name	Type	Description
capacity_shared	boolean	Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false.
max_throughput_iops	integer	Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced.
max_throughput_mbps	integer	Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced.

Name	Type	Description
min_throughput_iops	integer	Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
min_throughput_mbps	integer	Minimum throughput defined by this policy. It is specified in terms of Mbps. 0 means no minimum throughput is enforced.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos_policy

Name	Type	Description
_links	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
svm	svm	
uuid	string	

Create a QoS policy

POST /storage/qos/policies

Introduced In: 9.6

Creates a QoS policy.

Required properties

- `svm.uuid` or `svm.name` - The existing SVM owning the QoS policy.
- `name` - The name of the QoS policy.
- `fixed.*` or `adaptive.*` - Either of the fixed or adaptive parameters.

Default property values

- If `fixed.*` parameters are specified, then `capacity.shared` is set to false by default.

Related ONTAP commands

- `qos policy-group create`
- `qos adaptive-policy-group create`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.

Name	Type	Description
object_count	integer	Number of objects attached to this policy.
svm	svm	
uuid	string	

Example request

```
{
  "name": "extreme",
  "object_count": 0,
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8454147	The maximum limit for QoS policies has been reached.
8454154	The name specified for creating conflicts with an existing QoS policy name.
8454260	Invalid value for maximum and minimum fields. Valid values for max_throughput_iops and max_throughput_mbps combination is for the ratio of max_throughput_mbps and max_throughput_iops to be within 1 to 4096.
8454273	Invalid value for an adaptive field. Value should be non-zero.
8454277	The name specified for creating an adaptive QoS policy conflicts with an existing fixed QoS policy name.
8454278	The name specified for creating a fixed QoS policy conflicts with an existing adaptive QoS policy name.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

Name	Type	Description
capacity_shared	boolean	Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false.

Name	Type	Description
max_throughput_iops	integer	Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced.
max_throughput_mbps	integer	Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced.
min_throughput_iops	integer	Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
min_throughput_mbps	integer	Minimum throughput defined by this policy. It is specified in terms of Mbps. 0 means no minimum throughput is enforced.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos_policy

Name	Type	Description
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.

Name	Type	Description
object_count	integer	Number of objects attached to this policy.
svm	svm	
uuid	string	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a QoS policy

```
DELETE /storage/qos/policies/{uuid}
```

Introduced In: 9.6

Deletes a QoS policy. All QoS workloads associated with the policy are removed.

Related ONTAP commands

- `qos policy-group delete`

- qos adaptive-policy-group delete

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none">• Introduced in: 9.8
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a QoS policy

GET /storage/qos/policies/{uuid}

Introduced In: 9.6

Retrieves a specific QoS policy.

Related ONTAP commands

- `qos policy-group show`
- `qos adaptive-policy-group show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none">• Introduced in: 9.8
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
svm	svm	
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "extreme",
  "object_count": 0,
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

Name	Type	Description
capacity_shared	boolean	Specifies whether the capacities are shared across all objects that use this QoS policy-group. Default is false.
max_throughput_iops	integer	Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced.
max_throughput_mbps	integer	Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced.
min_throughput_iops	integer	Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
min_throughput_mbps	integer	Minimum throughput defined by this policy. It is specified in terms of Mbps. 0 means no minimum throughput is enforced.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a QoS policy

PATCH /storage/qos/policies/{uuid}

Introduced In: 9.6

Update a specific QoS policy.

Related ONTAP commands

- qos policy-group modify
- qos adaptive-policy-group modify

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	<ul style="list-style-type: none"> • Introduced in: 9.8

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.

Name	Type	Description
object_count	integer	Number of objects attached to this policy.
svm	svm	
uuid	string	

Example request

```
{
  "fixed": {
    "capacity_shared": null
  },
  "name": "extreme",
  "object_count": 0,
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8454147	The maximum limit for QoS policies has been reached.
8454154	The name specified for creating conflicts with an existing QoS policy name.
8454260	Invalid value for maximum and minimum fields. Valid values for max_throughput_iops and max_throughput_mbps combination is for the ratio of max_throughput_mbps and max_throughput_iops to be within 1 to 4096.
8454273	Invalid value for an adaptive field. Value should be non-zero.
8454277	The name specified for creating an adaptive QoS policy conflicts with an existing fixed QoS policy name.
8454278	The name specified for creating a fixed QoS policy conflicts with an existing adaptive QoS policy name.
8454286	Modifications on these cluster scoped preset policies is prohibited.
8454327	The existing fixed QoS policy cannot be modified to an adaptive QoS policy.
8454328	The existing adaptive QoS policy cannot be modified to a fixed QoS policy.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

adaptive

Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.

Name	Type	Description
absolute_min_iops	integer	Specifies the absolute minimum IOPS that is used as an override when the expected_iops is less than this value. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
expected_iops	integer	Expected IOPS. Specifies the minimum expected IOPS per TB allocated based on the storage object allocated size. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
peak_iops	integer	Peak IOPS. Specifies the maximum possible IOPS per TB allocated based on the storage object allocated size or the storage object used size.

fixed

QoS policy-groups define a fixed service level objective (SLO) for a storage object.

Name	Type	Description
max_throughput_iops	integer	Maximum throughput defined by this policy. It is specified in terms of IOPS. 0 means no maximum throughput is enforced.

Name	Type	Description
max_throughput_mbps	integer	Maximum throughput defined by this policy. It is specified in terms of Mbps. 0 means no maximum throughput is enforced.
min_throughput_iops	integer	Minimum throughput defined by this policy. It is specified in terms of IOPS. 0 means no minimum throughput is enforced. These floors are not guaranteed on non-AFF platforms or when FabricPool tiering policies are set.
min_throughput_mbps	integer	Minimum throughput defined by this policy. It is specified in terms of Mbps. 0 means no minimum throughput is enforced.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

qos_policy

Name	Type	Description
adaptive	adaptive	Adaptive QoS policy-groups define measurable service level objectives (SLOs) that adjust based on the storage object used space and the storage object allocated space.
fixed	fixed	QoS policy-groups define a fixed service level objective (SLO) for a storage object.
name	string	Name of the QoS policy.
object_count	integer	Number of objects attached to this policy.
svm	svm	

Name	Type	Description
uuid	string	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage storage qtrees

Storage qtrees endpoint overview

Overview

A qtree is a logically defined file system that can exist as a special subdirectory of the root directory within a FlexVol volume or a FlexGroup volume.

Qtree QoS Policy

Qtree QoS policy and settings enforce Service Level Objectives (SLO) on a qtree. SLO can be set by specifying "qos_policy.max_throughput_iops" and/or "qos_policy.max_throughput_mbps" or "qos_policy.min_throughput_iops".

Specifying "min_throughput_iops" is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying "qos_policy.name" or "qos_policy.uuid" properties.

Setting or assigning a QoS policy to a qtree is not supported if its containing volume or SVM already has a QoS policy attached, or a file or LUN in its containing volume already has a QoS policy attached.

Qtree APIs

The following APIs are used to create, retrieve, modify, and delete qtrees.

– POST /api/storage/qtrees

– GET /api/storage/qtrees

– GET /api/storage/qtrees/{volume-uuid}/{qtree-id}

– PATCH /api/storage/qtrees/{volume-uuid}/{qtree-id}

– DELETE /api/storage/qtrees/{volume-uuid}/{qtree-id}

Examples

Creating a qtree inside a volume for an SVM

This API is used to create a qtree inside a volume for an SVM.

The following example shows how to create a qtree in a FlexVol volume with a given security style, UNIX permissions, an export policy, and a QoS policy.

```
# The API:
POST /api/storage/qtrees

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/qtrees?return_records=true' -H
'accept: application/hal+json' -d @test_qtree_post.txt
test_qtree_post.txt (body):
{
  "svm": {
    "name": "svm1"
  },
  "volume": {
    "name": "fv"
  },
  "name": "qt1",
  "security_style": "unix",
  "unix_permissions": 744,
  "export_policy": {
    "name": "default"
  },
  "qos_policy": {
    "max_throughput_iops": 1000
  }
}
```

```

}
}

# The response:
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svml"
      },
      "volume": {
        "name": "fv"
      },
      "name": "qt1",
      "security_style": "unix",
      "unix_permissions": 744,
      "export_policy": {
        "name": "default"
      },
      "qos_policy": {
        "min_throughput_iops": 0,
        "max_throughput_iops": 1000,
        "max_throughput_mbps": 0,
        "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
        "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52"
      },
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/?volume.name=fv&name=qt1"
        }
      }
    },
    {
      "job": {
        "uuid": "84edef3c-4f6d-11e9-9a71-005056a7f717",
        "_links": {
          "self": {
            "href": "/api/cluster/jobs/84edef3c-4f6d-11e9-9a71-005056a7f717"
          }
        }
      }
    }
  ]
}

```

Retrieving qtrees

This API is used to retrieve qtrees.

The following example shows how to retrieve qtrees belonging to SVM *svm1* and volume *fv*. The `svm.name` and `volume.name` query parameters are used to find the required qtrees.

```
# The API:
GET /api/storage/qtrees

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/qtrees/?svm.name=svm1&volume.name=fv" -H 'accept:
application/hal+json'

# The response
{
  "records": [
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
          }
        }
      },
      "id": 0,
      "name": "",
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/0"
        }
      }
    }
  ]
}
```

```

    }
  },
  {
    "svm": {
      "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
      }
    },
    "volume": {
      "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
        }
      }
    },
    "id": 1,
    "name": "qt1",
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-
005056a7f717/1"
      }
    }
  },
  {
    "svm": {
      "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
      }
    },
    "volume": {
      "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
      "name": "fv",
      "_links": {
        "self": {

```

```

        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
      }
    },
    "id": 2,
    "name": "qt2",
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2"
      }
    }
  ],
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/?svm.name=svm1&volume.name=fv"
    }
  }
}

```

Retrieving properties of a specific qtree using a qtree identifier

This API is used to retrieve properties of a specific qtree using qtree.id.

The following example shows how to use the qtree identifier to retrieve all properties of the qtree using the `fields` query parameter.

```

# The API:
GET /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2?fields=*' -H 'accept: application/hal+json'
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  }
}

```



```

    }
  }
},
"volume": {
  "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
  "name": "fv",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
    }
  }
},
"id": 2,
"name": "qt2",
"security_style": "unix",
"unix_permissions": 744,
"export_policy": {
  "name": "default",
  "id": 12884901889,
  "_links": {
    "self": {
      "href": "/api/protocols/nfs/export-policies/12884901889"
    }
  }
},
"qos_policy": {
  "min_throughput_iops": 0,
  "max_throughput_iops": 1000,
  "max_throughput_mbps": 0,
  "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
  "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52",
  "_links": {
    "self": {
      "href": "/api/storage/qos/policies/39ac471f-ff35-11e9-b0f9-005056a7ab52"
    }
  }
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "iops_raw": {
    "read": 0,
    "write": 0,
    "other": 3,

```

```

    "total": 3
  },
  "throughput_raw": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"path": "/fv/qt2",
"_links": {
  "self": {
    "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2"
  }
}
}

```

Retrieving properties of a specific qtree using the qtree name

This API is used to retrieve properties of a specific qtree using `qtree.name`.

The following example shows how to retrieve all of the properties belonging to qtree `qt2`. The `svm.name` and `volume.name` query parameters are used here along with the qtree name.

```

# The API:
GET /api/storage/qtrees/

# The call:
curl -X GET 'https://<mgmt-
ip>/api/storage/qtrees/?svm.name=svm1&volume.name=fv&name=qt2&fields=*' -H
'accept: application/hal+json'
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",

```

```

    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
      }
    },
    "id": 2,
    "name": "qt2",
    "security_style": "unix",
    "unix_permissions": 744,
    "export_policy": {
      "name": "default",
      "id": 12884901889,
      "_links": {
        "self": {
          "href": "/api/protocols/nfs/export-policies/12884901889"
        }
      }
    },
    "qos_policy": {
      "min_throughput_iops": 0,
      "max_throughput_iops": 1000,
      "max_throughput_mbps": 0,
      "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab52",
      "name": "vs0_auto_gen_policy_39a9522f_ff35_11e9_b0f9_005056a7ab52",
      "_links": {
        "self": {
          "href": "/api/storage/qos/policies/39ac471f-ff35-11e9-b0f9-005056a7ab52"
        }
      }
    },
    "statistics": {
      "timestamp": "2019-04-09T05:50:42Z",
      "status": "ok",
      "iops_raw": {
        "read": 0,
        "write": 0,
        "other": 3,
        "total": 3
      },
      "throughput_raw": {
        "read": 0,
        "write": 0,
        "other": 0,

```

```
    "total": 0
  }
},
"_links": {
  "self": {
    "href": "/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2"
  }
}
}
```

Updating a qtree

This API is used to update a qtree.

The following example shows how to update properties in a qtree.

```
# The API:
PATCH /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2' -H 'accept: application/hal+json' -d
'@test_qtree_patch.txt'
test_qtree_patch.txt (body):
{
  "security_style": "mixed",
  "unix_permissions": 777,
  "export_policy": {
    "id": "9",
    "name": "exp1"
  },
  "qos_policy": {
    "uuid": "39ac471f-ff35-11e9-b0f9-005056a7ab53"
  }
}
```

Renaming a qtree

This API is used to rename a qtree.

The following example below shows how to rename a qtree with a new name.

```
# The API:
PATCH /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X PATCH 'https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/1' -H 'accept: application/hal+json' -d '{ "name": "new_qt1" }'
```

Deleting a qtree inside a volume of an SVM

This API is used to delete a qtree inside a volume of an SVM.

The following example shows how to delete a qtree.

```
# The API:
DELETE /api/storage/qtrees/{volume.uuid}/{id}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/qtrees/cb20da45-4f6b-11e9-9a71-005056a7f717/2" -H 'accept: application/hal+json'
```

Retrieve qtrees

GET /storage/qtrees

Introduced In: 9.6

Retrieves qtrees configured for all FlexVol volumes or FlexGroup volumes.

Use the `fields` query parameter to retrieve all properties of the qtree. If the `fields` query parameter is not used, then GET returns the qtree `name` and qtree `id` only.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `statistics.*`

Related ONTAP commands

- `qtree show`

Parameters

Name	Type	In	Required	Description
export_policy.id	integer	query	False	Filter by export_policy.id
export_policy.name	string	query	False	Filter by export_policy.name
name	string	query	False	Filter by name
path	string	query	False	Filter by path
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
qos_policy.max_throughput_mbps	integer	query	False	Filter by qos_policy.max_throughput_mbps <ul style="list-style-type: none">• Introduced in: 9.8
qos_policy.max_throughput_iops	integer	query	False	Filter by qos_policy.max_throughput_iops <ul style="list-style-type: none">• Introduced in: 9.8
qos_policy.uuid	string	query	False	Filter by qos_policy.uuid <ul style="list-style-type: none">• Introduced in: 9.8
qos_policy.min_throughput_iops	integer	query	False	Filter by qos_policy.min_throughput_iops <ul style="list-style-type: none">• Introduced in: 9.8

Name	Type	In	Required	Description
qos_policy.min_throughput_mbps	integer	query	False	Filter by qos_policy.min_throughput_mbps • Introduced in: 9.8
qos_policy.name	string	query	False	Filter by qos_policy.name • Introduced in: 9.8
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
unix_permissions	integer	query	False	Filter by unix_permissions
statistics.iops_raw.total	integer	query	False	Filter by statistics.iops_raw.total • Introduced in: 9.8
statistics.iops_raw.read	integer	query	False	Filter by statistics.iops_raw.read • Introduced in: 9.8
statistics.iops_raw.other	integer	query	False	Filter by statistics.iops_raw.other • Introduced in: 9.8
statistics.iops_raw.write	integer	query	False	Filter by statistics.iops_raw.write • Introduced in: 9.8

Name	Type	In	Required	Description
statistics.timestamp	string	query	False	Filter by statistics.timestamp • Introduced in: 9.8
statistics.status	string	query	False	Filter by statistics.status • Introduced in: 9.8
statistics.throughput_raw.total	integer	query	False	Filter by statistics.throughput_raw.total • Introduced in: 9.8
statistics.throughput_raw.read	integer	query	False	Filter by statistics.throughput_raw.read • Introduced in: 9.8
statistics.throughput_raw.other	integer	query	False	Filter by statistics.throughput_raw.other • Introduced in: 9.8
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write • Introduced in: 9.8
security_style	string	query	False	Filter by security_style
id	integer	query	False	Filter by id
fields	array[string]	query	False	Specify the fields to return.

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[qtree]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "export_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": "100",
        "name": "default"
      },
      "id": "1",
      "name": "string",
      "path": "/volume3/qtreen1",
      "qos_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "max_throughput_iops": "10000",
        "max_throughput_mbps": "500",
        "min_throughput_iops": "2000",
        "min_throughput_mbps": "500",
        "name": "performance",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "security_style": "string",
      "statistics": {
        "iops_raw": {
```

```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"unix_permissions": "0755",
"volume": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

qos_policy

Name	Type	Description
_links	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled with the next closest collection and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

qtree

A qtree is a directory at the top level of a volume to which a custom export policy (for fine-grained access control) and a quota rule can be applied, if required.

Name	Type	Description
_links	_links	
export_policy	export_policy	Export Policy
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
statistics	statistics	These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a qtree in a FlexVol or FlexGroup volume

POST /storage/qtrees

Introduced In: 9.6

Creates a qtree in a FlexVol or a FlexGroup volume.

After a qtree is created, the new qtree is assigned an identifier. This identifier is obtained using a qtree GET request. This identifier is used in the API path for the qtree PATCH and DELETE operations.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the qtree.
- `volume.uuid` or `volume.name` - Existing volume in which to create the qtree.
- `name` - Name for the qtree.

Recommended optional properties

If not specified in POST, the values are inherited from the volume.

- `security_style` - Security style for the qtree.
- `unix_permissions` - UNIX permissions for the qtree.
- `export_policy.name` or `export_policy.id` - Export policy of the SVM for the qtree.

Related ONTAP commands

- `qtree create`

Parameters

Name	Type	In	Required	Description
<code>return_timeout</code>	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
export_policy	export_policy	Export Policy
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

Example request

```
{
  "export_policy": {
    "id": "100",
    "name": "default"
  },
  "id": "1",
  "name": "string",
  "path": "/volume3/qtreet1",
  "qos_policy": {
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "security_style": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_permissions": "0755",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
917927	The specified volume was not found.
918232	Either <code>volume.name</code> or <code>volume.uuid</code> must be provided.
918236	The specified <code>volume.uuid</code> and <code>volume.name</code> refer to different volumes.
2621462	The specified SVM does not exist.
2621706	The specified <code>svm.uuid</code> and <code>svm.name</code> do not refer to the same SVM.
2621707	No SVM was specified. Either <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
5242886	Failed to create qtree.
5242951	Export Policy supplied does not belong to the specified Export Policy ID.
5242952	Export Policy ID specified is invalid.
5242953	Qtree name must be provided.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

export_policy

Export Policy

Name	Type	Description
id	integer	
name	string	

qos_policy

Name	Type	Description
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.

Name	Type	Description
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled with the next closest collection and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes does not have the latest data.

Name	Type	Description
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Required in POST

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

qtree

A qtree is a directory at the top level of a volume to which a custom export policy (for fine-grained access control) and a quota rule can be applied, if required.

Name	Type	Description
export_policy	export_policy	Export Policy

Name	Type	Description
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a qtree

`DELETE /storage/qtrees/{volume.uuid}/{id}`

Introduced In: 9.6

Deletes a qtree.

Related ONTAP commands

- `qtree delete`

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
id	string	path	True	Qtree ID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
918235	A volume with UUID was not found.
5242955	The UUID of the volume is required.
5242957	Failed to delete qtree with ID in volume and vserver.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve qtree properties

GET /storage/qtrees/{volume.uuid}/{id}

Introduced In: 9.6

Retrieves properties for a specific qtree identified by the `volume.uuid` and the `id` in the api path.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `statistics.*`

Related ONTAP commands

- `qtree show`

Parameters

Name	Type	In	Required	Description
<code>volume.uuid</code>	string	path	True	Volume UUID
<code>id</code>	string	path	True	Qtree ID
<code>fields</code>	array[string]	query	False	Specify the fields to return.

Response

```
Status: 200, Ok
```

Name	Type	Description
<code>_links</code>	_links	
<code>export_policy</code>	export_policy	Export Policy
<code>id</code>	integer	The identifier for the qtree, unique within the qtree's volume.
<code>name</code>	string	The name of the qtree. Required in POST; optional in PATCH.
<code>path</code>	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
<code>qos_policy</code>	qos_policy	

Name	Type	Description
security_style	string	Security style. Valid in POST or PATCH.
statistics	statistics	These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
svm	svm	Required in POST
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.
volume	volume	Required in POST

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "export_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": "100",
    "name": "default"
  },
  "id": "1",
  "name": "string",
  "path": "/volume3/qtreet1",
  "qos_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "security_style": "string",
  "statistics": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    }
  },
}
```

```

    "timestamp": "2017-01-25T11:20:13Z"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "unix_permissions": "0755",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
918235	A volume with UUID was not found.
5242956	Failed to obtain qtree with ID.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

qos_policy

Name	Type	Description
_links	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled with the next closest collection and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes does not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Required in POST

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
<code>_links</code>	_links	
<code>name</code>	string	The name of the volume.
<code>uuid</code>	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

error_arguments

Name	Type	Description
<code>code</code>	string	Argument code
<code>message</code>	string	Message argument

error

Name	Type	Description
<code>arguments</code>	array[error_arguments]	Message arguments
<code>code</code>	string	Error code
<code>message</code>	string	Error message
<code>target</code>	string	The target parameter that caused the error.

Update properties for a qtree

`PATCH /storage/qtrees/{volume.uuid}/{id}`

Introduced In: 9.6

Updates properties for a specific qtree.

Related ONTAP commands

- `qtree modify`

- `qtree rename`

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
id	string	path	True	Qtree ID
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Request Body

Name	Type	Description
export_policy	export_policy	Export Policy

Name	Type	Description
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.

Example request

```
{
  "export_policy": {
    "id": "100",
    "name": "default"
  },
  "id": "1",
  "name": "string",
  "path": "/volume3/qtree1",
  "qos_policy": {
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "security_style": "string",
  "unix_permissions": "0755"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
918235	A volume with UUID was not found.
5242951	Export policy supplied does not belong to the specified export policy ID.
5242955	The UUID of the volume is required.
5242956	Failed to obtain qtree with ID.
5242958	Failed to rename qtree in volume in SVM with ID.
5242959	Successfully renamed qtree but modify failed.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

export_policy

Export Policy

Name	Type	Description
id	integer	
name	string	

qos_policy

Name	Type	Description
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

iops_raw

The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput_raw

Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.

Name	Type	Description
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw IOPS and throughput performance numbers. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This should be used along with delta time to calculate the rate of I/O operations per unit of time.
status	string	Any errors associated with the sample. For example, if the aggregation of data over multiple nodes fails then any of the partial errors might be returned, "ok" on success, or "error" on any internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled with the next closest collection and tagged with "backfilled_data". "inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "negative_delta" is returned when an expected monotonically increasing value has decreased in value. "inconsistent_old_data" is returned when one or more nodes does not have the latest data.

Name	Type	Description
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This should be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

Required in POST

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Required in POST

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

qtree

A qtree is a directory at the top level of a volume to which a custom export policy (for fine-grained access control) and a quota rule can be applied, if required.

Name	Type	Description
export_policy	export_policy	Export Policy

Name	Type	Description
id	integer	The identifier for the qtree, unique within the qtree's volume.
name	string	The name of the qtree. Required in POST; optional in PATCH.
path	string	Client visible path to the qtree. This field is not available if the volume does not have a junction-path configured. Not valid in POST or PATCH.
qos_policy	qos_policy	
security_style	string	Security style. Valid in POST or PATCH.
unix_permissions	integer	The UNIX permissions for the qtree. Valid in POST or PATCH.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array [error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Storage quota

Storage quota endpoint overview

Overview

Quotas provide a way to restrict or track the files and space usage by a user, group, or qtree. Quotas are enabled for a specific FlexVol or a FlexGroup volume.

Quotas can have soft or hard limits. Soft limits cause ONTAP to send a notification when specified limits are exceeded. Hard limits prevent a write operation from succeeding when specified limits are exceeded.

Quota policy rule APIs

Quotas are defined as quota policy rules specific to FlexVol or FlexGroup volumes. Each quota rule has a type. The type can be either "user", "group", or "tree".

The following APIs can be used to perform create, retrieve, modify, and delete operations related to quota policy rules for a FlexVol or a FlexGroup volume.

– POST /api/storage/quota/rules

– GET /api/storage/quota/rules

– GET /api/storage/quota/rules/{rule-uuid}

– PATCH /api/storage/quota/rules/{rule-uuid}

– DELETE /api/storage/quota/rules/{rule-uuid}

Enabling and disabling quotas

After the quota rules are created, the quota rules need to be enabled on each FlexVol or FlexGroup volume for soft or hard limits to take effect in the filesystem. Enabling quotas can be done on a volume-by-volume basis.

The following APIs can be used to enable and disable and obtain the quota state for a FlexVol or a FlexGroup volume.

– PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"true"}

– PATCH /api/storage/volumes/{volume-uuid} -d '{"quota.enabled":"false"}

– GET /api/storage/volumes/{volume-uuid}/?fields=quota.state

Quota report APIs

Quota report records provide usage information for a user, group, or qtree against the quota limits configured on a FlexVol or a FlexGroup volume.

The following APIs can be used to retrieve quota reports associated with a FlexVol or a FlexGroup volume.

– GET /api/storage/quota/reports

– GET /api/storage/quota/reports/{volume-uuid}/{index}

Quota resize

Quota resize allows you to modify the quota limits directly in the filesystem.

It is important to note that quota must be enabled on a FlexVol or a FlexGroup volume for `quota resize` to take effect.

You can perform a `quota resize` using the quota policy rule PATCH API. If the quota is disabled on the volume, the quota policy rule PATCH API modifies the rule, but this does not affect the limits in the filesystem. The following API can be used to perform `quota resize` provided quota is enabled on the volume.

– PATCH /api/storage/quota/rules/{rule-uuid}
 The changed limits in the filesystem can be confirmed using the quota report REST API.

– GET /api/storage/quota/reports

Manage storage quota reports

Storage quota reports endpoint overview

Overview

Quota reports provide the current file and space consumption for a user, group, or qtree in a FlexVol or a FlexGroup volume.

Quota report APIs

The following APIs can be used to retrieve quota reports associated with a volume in ONTAP.

– GET /api/storage/quota/reports

– GET /api/storage/quota/reports/{volume_uuid}/{index}

Examples

Retrieving all the quota report records

This API is used to retrieve all the quota report records.

The following example shows how to retrieve quota report records for all FlexVol volumes and FlexGroup volumes.

```
# The API:
GET /api/storage/quota/reports

# The call:
```

```

curl -X GET "https://<mgmt-ip>/api/storage/quota/reports" -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
          }
        }
      },
      "index": 0,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/0"
        }
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",

```

```

    "name": "fg",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/314a328f-502d-11e9-8771-005056a7f717"
      }
    },
    "index": 1152921504606846976,
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-005056a7f717/1152921504606846976"
      }
    },
    {
      "svm": {
        "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
          }
        }
      },
      "volume": {
        "uuid": "314a328f-502d-11e9-8771-005056a7f717",
        "name": "fg",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/314a328f-502d-11e9-8771-005056a7f717"
          }
        }
      },
      "index": 3458764513820540928,
      "_links": {
        "self": {
          "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-005056a7f717/3458764513820540928"
        }
      },
      {
        "svm": {

```

```

    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "314a328f-502d-11e9-8771-005056a7f717",
    "name": "fg",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
      }
    }
  },
  "index": 4611686018427387904,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-
005056a7f717/4611686018427387904"
    }
  }
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "314a328f-502d-11e9-8771-005056a7f717",
    "name": "fg",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/314a328f-502d-11e9-8771-
005056a7f717"
      }
    }
  }
},

```

```

    "index": 5764607523034234880,
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/314a328f-502d-11e9-8771-005056a7f717/5764607523034234880"
      }
    }
  },
  {
    "svm": {
      "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
      "name": "svml",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
      }
    },
    "volume": {
      "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
        }
      }
    },
    "index": 0,
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/0"
      }
    }
  },
  {
    "svm": {
      "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
      "name": "svml",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
      }
    }
  },

```



```

    "volume": {
      "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
        }
      }
    },
    "index": 281474976710656,
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/281474976710656"
      }
    }
  },
  {
    "svm": {
      "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
      }
    },
    "volume": {
      "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
        }
      }
    },
    "index": 1152921504606846976,
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/1152921504606846976"
      }
    }
  },

```

```

{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
      }
    }
  },
  "index": 1153202979583557632,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/1153202979583557632"
    }
  }
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
      }
    }
  }
}

```

```

    }
  },
  "index": 2305843013508661248,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/2305843013508661248"
    }
  }
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
      }
    }
  },
  "index": 3458764513820540928,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/3458764513820540928"
    }
  }
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  }
}

```

```

    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
      }
    }
  },
  "index": 3459045988797251584,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/3459045988797251584"
    }
  }
},
{
  "svm": {
    "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
      }
    }
  },
  "volume": {
    "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717"
      }
    }
  },
  "index": 4611686018427387904,
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-005056a7f717/4611686018427387904"
    }
  }
}

```

```

    }
  },
  {
    "svm": {
      "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
      }
    },
    "volume": {
      "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-
005056a7f717"
        }
      }
    },
    "index": 4611967493404098560,
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/4611967493404098560"
      }
    }
  },
  {
    "svm": {
      "uuid": "b68f961b-4cee-11e9-930a-005056a7f717",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/b68f961b-4cee-11e9-930a-005056a7f717"
        }
      }
    },
    "volume": {
      "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-

```

```

005056a7f717"
    }
  }
},
"index": 5764607523034234880,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cb20da45-4f6b-11e9-9a71-
005056a7f717/5764607523034234880"
  }
}
}
],
"num_records": 15,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/"
  }
}
}
}

```

Retrieving a specific quota report record

This API is used to retrieve a specific quota report record.

The following example shows how to retrieve a single quota report user record.

```

# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-
11e9-8513-005056a7657c/281474976710656" -H 'accept: application/hal+json'

# Response for quota report user record:
{
  "svm": {
    "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
      }
    }
  }
}

```

```

    }
  },
  "volume": {
    "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
      }
    }
  },
  "index": 281474976710656,
  "type": "user",
  "users": [
    {
      "name": "fred",
      "id" : "300008"
    }
  ],
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
      }
    }
  },
  "space": {
    "hard_limit": 41943040,
    "soft_limit": 31457280,
    "used": {
      "total": 10567680,
      "soft_limit_percent": 34,
      "hard_limit_percent": 25
    }
  },
  "files": {
    "soft_limit": 30,
    "hard_limit": 40,
    "used": {
      "total": 11,
      "soft_limit_percent": 37,
      "hard_limit_percent": 28
    }
  },

```

```
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/281474976710656"
  }
}
}
```

Retrieving a single quota report multi-user record

```
# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/281474976710656" -H 'accept: application/hal+json'

# Response for quota report multi-user record:
{
  "svm": {
    "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
      }
    }
  },
  "volume": {
    "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
      }
    }
  },
  "index": 1153484454560268288,
  "type": "user",
  "users": [
    {
      "name": "fred",
```



```

    "id" : "300008"
  },
  {
    "name": "john",
    "id" : "300009"
  },
  {
    "name": "smith",
    "id" : "300010"
  }
],
"qtree": {
  "name": "qt1",
  "id": 1,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
    }
  }
},
"space": {
  "hard_limit": 41943040,
  "soft_limit": 31457280,
  "used": {
    "total": 10567680,
    "soft_limit_percent": 34,
    "hard_limit_percent": 25
  }
},
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,
    "hard_limit_percent": 28
  }
},
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/1153484454560268288"
  }
}
}

```

Retrieving a single quota report group record

```
# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/3459045988797251584" -H 'accept: application/hal+json'

# Response for quota report group record:
{
  "svm": {
    "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
      }
    }
  },
  "volume": {
    "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
      }
    }
  },
  "index": 3459045988797251584,
  "type": "group",
  "group": {
    "name" : "test_group",
    "id"   : "500009"
  },
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
      }
    }
  }
}
```

```

},
"space": {
  "hard_limit": 41943040,
  "soft_limit": 31457280,
  "used": {
    "total": 10567680,
    "soft_limit_percent": 34,
    "hard_limit_percent": 25
  }
},
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,
    "hard_limit_percent": 28
  }
},
"_links": {
  "self": {
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/3459045988797251584"
  }
}
}

```

Retrieving a single quota report tree record

```

# The API:
GET /api/storage/quota/reports/{volume.uuid}/{index}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/4612248968380809216" -H 'accept: application/hal+json'

# Response for quota report tree record:
{
  "svm": {
    "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
    "name": "svm1",

```

```

    "_links": {
      "self": {
        "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
      }
    },
    "volume": {
      "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
        }
      }
    },
    "index": 4612248968380809216,
    "type": "tree",
    "qtree": {
      "name": "qt1",
      "id": 1,
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
        }
      }
    },
    "space": {
      "hard_limit": 41943040,
      "soft_limit": 31457280,
      "used": {
        "total": 10567680,
        "soft_limit_percent": 34,
        "hard_limit_percent": 25
      }
    },
    "files": {
      "soft_limit": 30,
      "hard_limit": 40,
      "used": {
        "total": 11,
        "soft_limit_percent": 37,
        "hard_limit_percent": 28
      }
    },
    "_links": {
      "self": {

```

```
    "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/4612248968380809216"
  }
}
}
```

Retrieving only records enforced by non-default rules

```
# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports?show_default_records=false" -H 'accept: application/hal+json'

# Response from only non-default records
{
  "records": [
    {
      "svm": {
        "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
          }
        }
      },
      "volume": {
        "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
          }
        }
      },
      "index": 4612248968380809216,
      "type": "tree",
      "qtree": {
        "name": "qt1",
        "id": 1,
        "_links": {
```

```

    "self": {
      "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
    }
  },
  "space": {
    "hard_limit": 41943040,
    "soft_limit": 31457280,
    "used": {
      "total": 10567680,
      "soft_limit_percent": 34,
      "hard_limit_percent": 25
    }
  },
  "files": {
    "soft_limit": 30,
    "hard_limit": 40,
    "used": {
      "total": 11,
      "soft_limit_percent": 37,
      "hard_limit_percent": 28
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/4612248968380809216"
    }
  },
  {
    "svm": {
      "uuid": "5093e722-248e-11e9-96ee-005056a7657c",
      "name": "svml",
      "_links": {
        "self": {
          "href": "/api/svm/svms/5093e722-248e-11e9-96ee-005056a7657c"
        }
      }
    }
  },
  "volume": {
    "uuid": "cf480c37-2a6b-11e9-8513-005056a7657c",
    "name": "fv",
    "_links": {
      "self": {

```

```

    "href": "/api/storage/volumes/cf480c37-2a6b-11e9-8513-005056a7657c"
  }
},
"index": 1153484454560268288,
"type": "user",
"users": [
  {
    "name": "fred",
    "id" : "300008"
  },
  {
    "name": "john",
    "id" : "300009"
  },
  {
    "name": "smith",
    "id" : "300010"
  }
],
"qtree": {
  "name": "qt1",
  "id": 1,
  "_links": {
    "self": {
      "href": "/api/storage/qtrees/cf480c37-2a6b-11e9-8513-005056a7657c/1"
    }
  }
},
"space": {
  "hard_limit": 41943040,
  "soft_limit": 31457280,
  "used": {
    "total": 10567680,
    "soft_limit_percent": 34,
    "hard_limit_percent": 25
  }
},
"files": {
  "soft_limit": 30,
  "hard_limit": 40,
  "used": {
    "total": 11,
    "soft_limit_percent": 37,

```

```

        "hard_limit_percent": 28
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/cf480c37-2a6b-11e9-8513-005056a7657c/1153484454560268288"
      }
    }
  }
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports?show_default_records=false"
  }
}
}

```

Retrieving quota report records with query parameters

The following example shows how to retrieve tree type quota report records.

```

# The API:
GET /api/storage/quota/reports

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/quota/reports?type=tree" -H
'accept: application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
          }
        }
      }
    }
  ]
}

```



```

    },
    "volume": {
      "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
      "name": "fv",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-00505682cd5c"
        }
      }
    },
    "index": 2305843013508661248,
    "type": "tree",
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-00505682cd5c/2305843013508661248"
      }
    }
  },
  {
    "svm": {
      "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
      "name": "svm1",
      "_links": {
        "self": {
          "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
        }
      }
    },
    "volume": {
      "uuid": "a5ceebd2-6ccf-11ea-bc35-005056823577",
      "name": "fg",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/a5ceebd2-6ccf-11ea-bc35-005056823577"
        }
      }
    },
    "index": 2305843013508661248,
    "type": "tree",
    "_links": {
      "self": {
        "href": "/api/storage/quota/reports/a5ceebd2-6ccf-11ea-bc35-005056823577/2305843013508661248"
      }
    }
  }
]

```

```

    }
  }
},
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/storage/quota/reports?type=tree"
  }
}
}
}

```

Retrieving all the quota reports of a specific volume and the files fields

```

# The API:
GET /api/storage/quota/reports

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/quota/reports?volume.name=fv&fields=files" -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
        "name": "svm1",
        "_links": {
          "self": {
            "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
          }
        }
      },
      "volume": {
        "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
        "name": "fv",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-
00505682cd5c"

```

```

    }
  },
  "index": 410328290557952,
  "files": {
    "soft_limit": 20,
    "hard_limit": 30,
    "used": {
      "total": 0,
      "soft_limit_percent": 0,
      "hard_limit_percent": 0
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-00505682cd5c/410328290557952"
    }
  }
},
{
  "svm": {
    "uuid": "903e54ee-6ccf-11ea-bc35-005056823577",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/903e54ee-6ccf-11ea-bc35-005056823577"
      }
    }
  },
  "volume": {
    "uuid": "8812b000-6e1e-11ea-9bad-00505682cd5c",
    "name": "fv",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/8812b000-6e1e-11ea-9bad-00505682cd5c"
      }
    }
  },
  "index": 2305843013508661248,
  "files": {
    "soft_limit": 200,
    "hard_limit": 400,
    "used": {
      "total": 4,

```

```

        "soft_limit_percent": 2,
        "hard_limit_percent": 1
    }
},
"_links": {
    "self": {
        "href": "/api/storage/quota/reports/8812b000-6e1e-11ea-9bad-00505682cd5c/2305843013508661248"
    }
}
],
"num_records": 2,
"_links": {
    "self": {
        "href": "/api/storage/quota/reports?volume.name=fv&fields=files"
    }
}
}
}

```

Retrieve the quota report records for all FlexVol and FlexGroup volumes

GET /storage/quota/reports

Introduced In: 9.6

Retrieves the quota report records for all FlexVol volumes and FlexGroup volumes.

Related ONTAP commands

- `quota report`

Parameters

Name	Type	In	Required	Description
users.id	string	query	False	Filter by users.id
users.name	string	query	False	Filter by users.name
space.used.soft_limit_percent	integer	query	False	Filter by space.used.soft_limit_percent
space.used.total	integer	query	False	Filter by space.used.total

Name	Type	In	Required	Description
space.used.hard_limit_percent	integer	query	False	Filter by space.used.hard_limit_percent
space.hard_limit	integer	query	False	Filter by space.hard_limit
space.soft_limit	integer	query	False	Filter by space.soft_limit
type	string	query	False	Filter by type
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
group.name	string	query	False	Filter by group.name
group.id	string	query	False	Filter by group.id
qtree.id	integer	query	False	Filter by qtree.id
qtree.name	string	query	False	Filter by qtree.name
files.soft_limit	integer	query	False	Filter by files.soft_limit
files.hard_limit	integer	query	False	Filter by files.hard_limit
files.used.soft_limit_percent	integer	query	False	Filter by files.used.soft_limit_percent
files.used.total	integer	query	False	Filter by files.used.total
files.used.hard_limit_percent	integer	query	False	Filter by files.used.hard_limit_percent
index	integer	query	False	Filter by index
specifier	string	query	False	Filter by specifier

Name	Type	In	Required	Description
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
show_default_records	boolean	query	False	<p>The default is true for GET calls. When set to false, the default records are not reported.</p> <ul style="list-style-type: none"> • Introduced in: 9.7 • Default value: 1
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[quota_report]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "files": {
        "hard_limit": 0,
        "soft_limit": 0,
        "used": {
          "hard_limit_percent": 0,
          "soft_limit_percent": 0,
          "total": 0
        }
      },
      "group": {
        "id": "string",
        "name": "string"
      },
      "index": 0,
      "qtree": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": "1",
        "name": "qt1"
      },
      "space": {
        "hard_limit": 0,
        "soft_limit": 0,
        "used": {
          "hard_limit_percent": 0,
```



```

        "soft_limit_percent": 0,
        "total": 0
    }
},
"specifier": "string",
"svm": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"type": "string",
"users": [
    {
        "id": "string",
        "name": "string"
    }
],
"volume": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

used

Name	Type	Description
hard_limit_percent	integer	Total files used as a percentage of file hard limit
soft_limit_percent	integer	Total files used as a percentage of file soft limit
total	integer	Total files used

files

Name	Type	Description
hard_limit	integer	File hard limit
soft_limit	integer	File soft limit
used	used	

group

This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

used

Name	Type	Description
hard_limit_percent	integer	Total space used as a percentage of space hard limit
soft_limit_percent	integer	Total space used as a percentage of space soft limit
total	integer	Total space used

space

Name	Type	Description
hard_limit	integer	Space hard limit in bytes
soft_limit	integer	Space soft limit in bytes
used	used	

svm

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

quota_report

Name	Type	Description
_links	_links	
files	files	
group	group	<p>This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifier.</p>

Name	Type	Description
index	integer	Index that identifies a unique quota record. Valid in URL.
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
specifier	string	Quota specifier
svm	svm	
type	string	Quota type associated with the quota record.
users	array[users]	This parameter specifies the target user or users associated with the given quota report record. This parameter is available for user quota records and is not available for group or tree quota records. The target user or users are identified by a user name and user identifier. The user name can be a UNIX user name or a Windows user name, and the identifier can be a UNIX user identifier or a Windows security identifier.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a specific quota report record

GET /storage/quota/reports/{volume.uuid}/{index}

Introduced In: 9.6

Retrieves a specific quota report record.

Related ONTAP commands

- `quota report`

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
index	integer	path	True	Quota report index
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
files	files	
group	group	This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifier.
index	integer	Index that identifies a unique quota record. Valid in URL.
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
specifier	string	Quota specifier
svm	svm	
type	string	Quota type associated with the quota record.

Name	Type	Description
users	array[users]	This parameter specifies the target user or users associated with the given quota report record. This parameter is available for user quota records and is not available for group or tree quota records. The target user or users are identified by a user name and user identifier. The user name can be a UNIX user name or a Windows user name, and the identifier can be a UNIX user identifier or a Windows security identifier.
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "files": {
    "hard_limit": 0,
    "soft_limit": 0,
    "used": {
      "hard_limit_percent": 0,
      "soft_limit_percent": 0,
      "total": 0
    }
  },
  "group": {
    "id": "string",
    "name": "string"
  },
  "index": 0,
  "qtree": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": "1",
    "name": "qt1"
  },
  "space": {
    "hard_limit": 0,
    "soft_limit": 0,
    "used": {
      "hard_limit_percent": 0,
      "soft_limit_percent": 0,
      "total": 0
    }
  },
  "specifier": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    }
  }
}
```

```

    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
918235	A volume with UUID was not found.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

used

Name	Type	Description
hard_limit_percent	integer	Total files used as a percentage of file hard limit
soft_limit_percent	integer	Total files used as a percentage of file soft limit
total	integer	Total files used

files

Name	Type	Description
hard_limit	integer	File hard limit
soft_limit	integer	File soft limit
used	used	

group

This parameter specifies the target group associated with the given quota report record. This parameter is available for group quota records and is not available for user or tree quota records. The target group is identified by a UNIX group name and UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a

user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

used

Name	Type	Description
hard_limit_percent	integer	Total space used as a percentage of space hard limit
soft_limit_percent	integer	Total space used as a percentage of space soft limit
total	integer	Total space used

space

Name	Type	Description
hard_limit	integer	Space hard limit in bytes
soft_limit	integer	Space soft limit in bytes
used	used	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage storage quota policy rules

Storage quota rules endpoint overview

Overview

Quotas are defined in quota rules specific to FlexVol volumes or FlexGroup volumes. Each quota rule has a type. The type can be "user", "group", or "tree".

– User rules must have the user property and qtree property.

– Group rules must have the group property and qtree property.

– Tree rules must have the qtree property and not have the user or group property.

Quota policy rule APIs

The following APIs can be used to perform create, retrieve, modify, and delete operations related to quota policy rules.

– POST /api/storage/quota/rules

– GET /api/storage/quota/rules

– GET /api/storage/quota/rules/{rule-uuid}

– PATCH /api/storage/quota/rules/{rule-uuid}

– DELETE /api/storage/quota/rules/{rule-uuid}

Examples

Retrieving all quota policy rules

This API is used to retrieve all quota policy rules.

The following example shows how to retrieve quota policy rules for FlexVol volumes and FlexGroup volumes.

```
# The API:
GET /api/storage/quota/rules

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules' -H 'accept:
application/hal+json'

# The response:
{
  "records": [
    {
      "svm": {
```



```

    "uuid": "038545f8-9ff8-11e8-bce6-005056a73bed",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
      }
    }
  },
  "volume": {
    "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ab3df793-0f02-43c6-9514-4f142fc8cc92"
      }
    }
  },
  "uuid": "66319cbe-b837-11e8-9c5a-005056a7e88c",
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/66319cbe-b837-11e8-9c5a-005056a7e88c"
    }
  }
},
{
  "svm": {
    "uuid": "038545f8-9ff8-11e8-bce6-005056a73bed",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/038545f8-9ff8-11e8-bce6-005056a73bed"
      }
    }
  },
  "volume": {
    "uuid": "ab3df793-0f02-43c6-9514-4f142fc8cc92",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/ab3df793-0f02-43c6-9514-4f142fc8cc92"
      }
    }
  }
},

```

```

    "uuid": "dbd5b443-b7a4-11e8-bc58-005056a7e88c",
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/dbd5b443-b7a4-11e8-bc58-005056a7e88c"
      }
    }
  },
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules"
    }
  }
}

```

Retrieving a specific quota policy rule

This API is used to retrieve a quota policy rule for a specific qtree.

The following example shows how to retrieve a quota policy user rule for a specific qtree.

```

# The API:
GET /api/storage/quota/rules/{uuid}

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/264a9e0b-2e03-11e9-a610-005056a7b72d' -H 'accept: application/hal+json'

# Response for a user rule at a qtree level:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",

```

```

    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/clb64eea-ca8b-45ec-9397-
ab489830d268"
      }
    },
    "uuid": "264a9e0b-2e03-11e9-a610-005056a7b72d",
    "type": "user",
    "users": [ {"name" : "fred"} ],
    "qtree": {
      "name": "qt1",
      "id": 1,
      "_links": {
        "self": {
          "href": "/api/storage/qtrees/clb64eea-ca8b-45ec-9397-
ab489830d268/1"
        }
      }
    },
    "user_mapping": on,
    "space": {
      "hard_limit": 1222800,
      "soft_limit": 51200
    },
    "files": {
      "hard_limit": 100,
      "soft_limit": 80
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/264a9e0b-2e03-11e9-a610-
005056a7b72d"
      }
    }
  }
}

```

Retrieving a quota policy multi-user rule at the volume level

```

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/0ab84fba-19aa-11e9-

```

```
a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a multi-user rule at volume level:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svml",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "voll",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-ab489830d268"
      }
    }
  },
  "uuid": "0ab84fba-19aa-11e9-a04d-005056a72f42",
  "type": "user",
  "users": [
    {
      "name": "sam",
    },
    {
      "name": "smith",
    },
    {
      "id": "300010",
    },
  ],
  "space": {
    "hard_limit": 1222800,
    "soft_limit": 51200
  },
  "files": {
    "hard_limit": 100,
    "soft_limit": 80
  },
  "_links": {
    "self": {
```

```
    "href": "/api/storage/quota/rules/0ab84fba-19aa-11e9-a04d-  
005056a72f42"  
  }  
}  
}
```

Retrieving a quota policy default tree rule

```

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/4a276b8c-1753-11e9-8101-005056a760e0' -H 'accept: application/hal+json'

# Response for a default tree rule:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "voll",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-ab489830d268"
      }
    }
  },
  "uuid": "4a276b8c-1753-11e9-8101-005056a760e0",
  "type": "tree",
  "qtree": {
    "name": ""
  },
  "space": {
    "hard_limit": 1034000,
    "soft_limit": 51200
  },
  "files": {
    "hard_limit": 20,
    "soft_limit": 10
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/4a276b8c-1753-11e9-8101-005056a760e0"
    }
  }
}

```

Retrieving a quota policy tree rule for a specific qtree

```
# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/49b1134f-19ab-11e9-a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a tree rule for a specific qtree:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-ab489830d268"
      }
    }
  },
  "uuid": "49b1134f-19ab-11e9-a04d-005056a72f42",
  "type": "tree",
  "qtree": {
    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/c1b64eea-ca8b-45ec-9397-ab489830d268/1"
      }
    }
  },
  "space": {
    "hard_limit": 1048576,
    "soft_limit": 838861
  },
  "files": {
```

```

    "hard_limit": 100,
    "soft_limit": 40
  },
  "_links": {
    "self": {
      "href": "/api/storage/quota/rules/49b1134f-19ab-11e9-a04d-005056a72f42"
    }
  }
}

```

Retrieving a quota policy group rule for a specific qtree

```

# The call:
curl -X GET 'https://<mgmt-ip>/api/storage/quota/rules/b9236852-19ab-11e9-a04d-005056a72f42' -H 'accept: application/hal+json'

# Response for a group rule:
{
  "svm": {
    "uuid": "fd5db15a-15b9-11e9-a6ad-005056a760e0",
    "name": "svm1",
    "_links": {
      "self": {
        "href": "/api/svm/svms/fd5db15a-15b9-11e9-a6ad-005056a760e0"
      }
    }
  },
  "volume": {
    "uuid": "c1b64eea-ca8b-45ec-9397-ab489830d268",
    "name": "vol1",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/c1b64eea-ca8b-45ec-9397-ab489830d268"
      }
    }
  },
  "uuid": "b9236852-19ab-11e9-a04d-005056a72f42",
  "type": "group",
  "group": {"name" : "group1"},
  "qtree": {

```



```

    "name": "qt1",
    "id": 1,
    "_links": {
      "self": {
        "href": "/api/storage/qtrees/clb64eea-ca8b-45ec-9397-
ab489830d268/1"
      }
    },
    "space": {
      "hard_limit": 2097152,
      "soft_limit": 1572864
    },
    "files": {
      "hard_limit": 250,
      "soft_limit": 200
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/b9236852-19ab-11e9-a04d-
005056a72f42"
      }
    }
  }
}

```

Creating a quota policy rule

This API is used to create a new quota policy rule. When an explicit rule or a qtree-scoped rule of a type is created on a volume, a default rule of the same type is automatically added if it does not already exist on the volume.

The following example shows how to create a quota policy user rule using POST.

```

# The API:
POST /api/storage/quota/rules

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/quota/rules?return_records=true' -H 'accept:
application/hal+json' -d @test_quota_post.txt
test_quota_post.txt (body):
{
  "svm": {

```

```

    "name": "svm1"
  },
  "volume": {
    "name": "vol1"
  },
  "type": "user",
  "users": [ {"name" : "jsmith"} ],
  "qtree": {
    "name": "qt1"
  },
  "user_mapping": "on",
  "space": {
    "hard_limit": 8192,
    "soft_limit": 1024
  },
  "files": {
    "hard_limit": 20,
    "soft_limit": 10
  }
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "3220eea6-5049-11e9-bfb7-005056a7f717",
      "type": "user",
      "users": [
        {
          "name" : "jsmith"
        }
      ],
      "qtree": {
        "name": "qt1"
      },
      "user_mapping": "on",
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      }
    }
  ]
}

```

```

    },
    "files": {
      "hard_limit": 20,
      "soft_limit": 10
    },
    "_links": {
      "self": {
        "href": "/api/storage/quota/rules/3220eea6-5049-11e9-bfb7-005056a7f717"
      }
    }
  },
  "job": {
    "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
      }
    }
  }
}

```

Creating a quota policy group rule using POST.

```

# The API:
POST /api/storage/quota/rules

# The call:
curl -X POST 'https://<mgmt-ip>/api/storage/quota/rules?return_records=true' -H 'accept: application/hal+json' -d @test_quota_post.txt
test_quota_post.txt (body):
{
  "svm": {
    "name": "svm1"
  },
  "volume": {
    "name": "vol1"
  },
  "type": "group",
  "group": {

```

```

    "name" : "test_group1"
  },
  "qtree": {
    "name": "qt1"
  },
  "space": {
    "hard_limit": 8192,
    "soft_limit": 1024
  },
  "files": {
    "hard_limit": 20,
    "soft_limit": 10
  }
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svm1"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "3b130f7d-504a-11e9-bfb7-005056a7f717",
      "type": "group",
      "group": {
        "name" : "test_group1"
      },
      "qtree": {
        "name": "qt1"
      },
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
      "files": {
        "hard_limit": 20,
        "soft_limit": 10
      },
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/3b130f7d-504a-11e9-bfb7-005056a7f717"
        }
      }
    }
  ]
}

```

```

    }
  }
},
"job": {
  "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
  "_links": {
    "self": {
      "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
    }
  }
}
}
}

```

Creating a quota policy tree rule using POST

```

# The API:
POST /api/storage/quota/rules

# The call:
curl -X POST 'https://<mgmt-
ip>/api/storage/quota/rules?return_records=true' -H 'accept:
application/hal+json' -d @test_quota_post.txt
test_quota_post.txt(body):
{
  "svm": {
    "name": "svm1"
  },
  "volume": {
    "name": "vol1"
  },
  "type": "tree",
  "qtree": {
    "name": "qt1"
  },
  "space": {
    "hard_limit": 8192,
    "soft_limit": 1024
  },
  "files": {
    "hard_limit": 20,
    "soft_limit": 10
  }
}

```

```

}
}

# The response
{
  "num_records": 1,
  "records": [
    {
      "svm": {
        "name": "svml"
      },
      "volume": {
        "name": "fv"
      },
      "uuid": "e5eb03be-504a-11e9-bfb7-005056a7f717",
      "type": "tree",
      "qtree": {
        "name": "qt1"
      },
      "space": {
        "hard_limit": 8192,
        "soft_limit": 1024
      },
      "files": {
        "hard_limit": 20,
        "soft_limit": 10
      },
      "_links": {
        "self": {
          "href": "/api/storage/quota/rules/e5eb03be-504a-11e9-bfb7-005056a7f717"
        }
      }
    },
    {
      "job": {
        "uuid": "32223924-5049-11e9-bfb7-005056a7f717",
        "_links": {
          "self": {
            "href": "/api/cluster/jobs/32223924-5049-11e9-bfb7-005056a7f717"
          }
        }
      }
    }
  ]
}

```

Updating the quota policy rule

This API is used to update a quota policy rule.

The following example shows how to update a quota policy rule.

```
# The API:
PATCH /storage/quota/rules/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/quota/rules/364d38eb-8e87-11e8-a806-005056a7e73a" -H 'accept: application/hal+json' -d
"@test_quota_patch.txt"
test_quota_patch.txt (body):
{
  "space": {
    "hard_limit": 16554,
    "soft_limit": 8192
  },
  "files": {
    "hard_limit": 40,
    "soft_limit": 20
  }
}
```

Deleting the quota policy rule

This API is used to delete a quota policy rule.

The following example shows how to delete a quota policy rule.

```
# The API:
DELETE /storage/quota/rules/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/quota/rules/364d38eb-8e87-11e8-a806-005056a7e73a" -H 'accept: application/hal+json'
```

Retrieve quota policy rules for all FlexVol and FlexGroup volumes

GET /storage/quota/rules

Introduced In: 9.6

Retrieves quota policy rules configured for all FlexVol volumes and FlexGroup volumes.

Related ONTAP commands

- `quota policy rule show`

Parameters

Name	Type	In	Required	Description
type	string	query	False	Filter by type
users.name	string	query	False	Filter by users.name
users.id	string	query	False	Filter by users.id
space.soft_limit	integer	query	False	Filter by space.soft_limit
space.hard_limit	integer	query	False	Filter by space.hard_limit
group.name	string	query	False	Filter by group.name
group.id	string	query	False	Filter by group.id
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
user_mapping	boolean	query	False	Filter by user_mapping
uuid	string	query	False	Filter by uuid
files.soft_limit	integer	query	False	Filter by files.soft_limit
files.hard_limit	integer	query	False	Filter by files.hard_limit
qtree.id	integer	query	False	Filter by qtree.id

Name	Type	In	Required	Description
qtree.name	string	query	False	Filter by qtree.name
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[quota_rule]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "group": {
        "id": "string",
        "name": "string"
      },
      "qtree": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "id": "1",
        "name": "qt1"
      },
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "type": "string",
      "users": [
        {
          "id": "string",
          "name": "string"
        }
      ]
    }
  ]
}
```

```

    ],
    "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
    "volume": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
  }
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

files

Name	Type	Description
hard_limit	integer	This parameter specifies the hard limit for files. This is valid in POST or PATCH.
soft_limit	integer	This parameter specifies the soft limit for files. This is valid in POST or PATCH.

group

This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.
soft_limit	integer	This parameter specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

quota_rule

Name	Type	Description
_links	_links	
files	files	

Name	Type	Description
group	group	This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	
type	string	This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.

Name	Type	Description
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.
users	array[users]	This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a quota policy rule for a FlexVol or a FlexGroup volume

POST /storage/quota/rules

Introduced In: 9.6

Creates a quota policy rule for a FlexVol or a FlexGroup volume.

Important notes:

- Unlike CLI/ONTAPI, the `quota_policy` input is not needed for POST.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the qtree.
- `volume.uuid` or `volume.name` - Existing volume in which to create the qtree.
- `type` - Quota type for the rule. This type can be `user`, `group`, or `tree`.
- `users.name` or `user.id` - If the quota type is `user`, this property takes the user name or user ID. For default user quota rules, the user name must be specified as `""`.
- `group.name` or `group.id` - If the quota type is `group`, this property takes the group name or group ID. For default group quota rules, the group name must be specified as `""`.
- `qtree.name` - Qtree for which to create the rule. For default tree rules, the qtree name must be specified as `""`.

Recommended optional properties

- `space.hard_limit` - Specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes.

- `space.soft_limit` - Specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes.
- `files.hard_limit` - Specifies the hard limit for files.
- `files.hard_limit` - Specifies the soft limit for files.
- `user_mapping` - Specifies the `user_mapping`. This property is valid only for quota policy rules of type `user`.

Related ONTAP commands

- `quota policy rule create`

Parameters

Name	Type	In	Required	Description
<code>return_timeout</code>	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
files	files	
group	group	<p>This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.</p>
qtree	qtree	<p>This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.</p>

Name	Type	Description
space	space	
svm	svm	
type	string	This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.
users	array[users]	This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.

Name	Type	Description
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

Example request

```
{
  "group": {
    "id": "string",
    "name": "string"
  },
  "qtree": {
    "id": "1",
    "name": "qt1"
  },
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
917927	The specified volume was not found.
918232	Either <code>volume.name</code> or <code>volume.uuid</code> must be provided.
918236	The specified <code>volume.uuid</code> and <code>volume.name</code> refer to different volumes.
2621462	The specified SVM does not exist.
2621706	The specified <code>svm.uuid</code> and <code>svm.name</code> do not refer to the same SVM.
2621707	No SVM was specified. Either <code>svm.name</code> or <code>svm.uuid</code> must be supplied.
5308501	Mapping from Windows user to UNIX user for user rule was unsuccessful.
5308502	Mapping from UNIX user to Windows user for user rule was unsuccessful.
5308552	Failed to get default quota policy name for SVM.
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
5308562	<code>users</code> is a required input for creating a user rule and <code>group</code> is not allowed.
5308563	<code>group</code> is a required input for creating a group rule and <code>users</code> is not allowed.

Error Code	Description
5308564	qtree.name is a required input for creating a tree rule and users and group are not allowed.
5308565	Only one of name or id is allowed for each entry in the users array.
5308566	Only one of name or id is allowed for group.
5308568	Quota policy rule create operation succeeded, but quota resize failed due to internal error. To activate the rule, disable and enable quotas for this volume.
5308571	Quota policy rule create operation succeeded, but quota resize is skipped. To activate the rule, disable and enable quotas for this volume.
5308573	Input value is greater than limit for field.
5308574	Input value is out of range for field.
5308575	Input value is incorrectly larger than listed field.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

files

Name	Type	Description
hard_limit	integer	This parameter specifies the hard limit for files. This is valid in POST or PATCH.
soft_limit	integer	This parameter specifies the soft limit for files. This is valid in POST or PATCH.

group

This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
id	integer	The unique identifier for a qtree.

Name	Type	Description
name	string	The name of the qtree.

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.
soft_limit	integer	This parameter specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

quota_rule

Name	Type	Description
files	files	
group	group	<p>This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.</p>

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	
type	string	This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.

Name	Type	Description
users	array[users]	This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a quota policy rule

DELETE /storage/quota/rules/{uuid}

Introduced In: 9.7

Deletes a quota policy rule.

Related ONTAP commands

- `quota policy rule delete`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Rule UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5308545	The specified quota rule UUID is invalid.
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
5308569	Quota policy rule delete operation succeeded, but quota resize failed due to internal error.
5308572	Quota policy rule delete operation succeeded, however the rule is still being enforced. To stop enforcing the rule, disable quotas and enable them again for this volume.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve quota policy rule properties

GET /storage/quota/rules/{uuid}

Introduced In: 9.7

Retrieves properties for a specific quota policy rule.

Related ONTAP commands

- `quota policy rule show`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Rule UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
files	files	
group	group	This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
qtree	qtree	This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.
space	space	
svm	svm	
type	string	This parameter specifies the quota policy rule type. This is required in POST only and can take either one of the "user", "group" or "tree" values.
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.

Name	Type	Description
users	array[users]	This parameter specifies the target user to which the user quota policy rule applies. This parameter takes single or multiple user names or identifiers. This parameter is valid only for the POST operation of a user quota policy rule. If this parameter is used as an input to create a group or a tree quota policy rule, the POST operation will fail with an appropriate error. For POST, this input parameter takes either a user name or a user identifier, not both. For default quota rules, the user name must be chosen and specified as "". For explicit user quota rules, this parameter can indicate either a user name or user identifier. The user name can be a UNIX user name or a Windows user name. If a name contains a space, enclose the entire value in quotes. A UNIX user name cannot include a backslash () or an @ sign; user names with these characters are treated as Windows names. The user identifier can be a UNIX user identifier or a Windows security identifier. For multi-user quota, this parameter can contain multiple user targets separated by a comma.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "group": {
    "id": "string",
    "name": "string"
  },
  "qtrees": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "id": "1",
    "name": "qt1"
  },
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "users": [
    {
      "id": "string",
      "name": "string"
    }
  ],
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

```
}  
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5308544	The specified quota rule UUID is invalid.
5308545	Unable to retrieve rule for the specified quota rule UUID.
5308576	Parameter <code>show_default_records</code> only allowed for GET collection.

Name	Type	Description
error	error	

Example error

```
{  
  "error": {  
    "arguments": [  
      {  
        "code": "string",  
        "message": "string"  
      }  
    ],  
    "code": "4",  
    "message": "entry doesn't exist",  
    "target": "uuid"  
  }  
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

files

Name	Type	Description
hard_limit	integer	This parameter specifies the hard limit for files. This is valid in POST or PATCH.
soft_limit	integer	This parameter specifies the soft limit for files. This is valid in POST or PATCH.

group

This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
_links	_links	
id	integer	The unique identifier for a qtree.
name	string	The name of the qtree.

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.
soft_limit	integer	This parameter specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update quota policy rule properties

PATCH /storage/quota/rules/{uuid}

Introduced In: 9.7

Updates properties of a specific quota policy rule.

Important notes:

- The quota resize functionality is supported with the PATCH operation.
- Quota resize allows you to modify the quota limits, directly in the filesystem.
- The quota must be enabled on a FlexVol or a FlexGroup volume for `quota resize` to take effect.
- If the quota is disabled on the volume, the quota policy rule PATCH API modifies the rule, but this does not affect the limits in the filesystem.

Related ONTAP commands

- `quota policy rule modify`
- `quota resize`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Rule UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
files	files	
space	space	
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.

Name	Type	Description
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.

Example request

```
{
  "uuid": "5f1d13a7-f401-11e8-ac1a-005056a7c3b9"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5308501	Mapping from Windows user to UNIX user for user rule was unsuccessful.
5308502	Mapping from UNIX user to Windows user for user rule was unsuccessful.
5308545	The specified quota rule UUID is invalid.

Error Code	Description
5308561	Failed to obtain volume quota state or invalid quota state obtained for volume.
5308567	Quota policy rule modify operation succeeded, but quota resize failed due to internal error.
5308573	Input value is greater than limit for field.
5308574	Input value is out of range for field.
5308575	Input value is incorrectly larger than listed field.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

files

Name	Type	Description
hard_limit	integer	This parameter specifies the hard limit for files. This is valid in POST or PATCH.
soft_limit	integer	This parameter specifies the soft limit for files. This is valid in POST or PATCH.

group

This parameter specifies the target group to which the group quota policy rule applies. This parameter takes a group name or identifier. This parameter is only valid for the POST operation of a group quota policy rule. The POST operation will fail with an appropriate error if this parameter is used as an input to create a user or a tree quota policy rule. This input parameter for POST takes either a group name or a group identifier, but not both. For default quota rules, the group name must be chosen and should be specified as "". For explicit group quota rules, this parameter can contain a UNIX group name or a UNIX group identifier.

Name	Type	Description
id	string	Quota target group ID
name	string	Quota target group name

qtree

This parameter specifies the target qtree to which the user/group/tree quota policy rule applies. For a user/group quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST. For a user/group quota policy rule at volume level, this parameter is not valid in GET or POST. For a tree quota policy rule, this parameter is mandatory and is valid in both POST and GET. For a default tree quota policy rule, this parameter needs to be specified as "". For a tree quota policy rule at qtree level, this parameter takes a qtree name and is valid in GET or POST.

Name	Type	Description
id	integer	The unique identifier for a qtree.

Name	Type	Description
name	string	The name of the qtree.

space

Name	Type	Description
hard_limit	integer	This parameter specifies the space hard limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.
soft_limit	integer	This parameter specifies the space soft limit, in bytes. If less than 1024 bytes, the value is rounded up to 1024 bytes. Valid in POST or PATCH. For a POST operation where the parameter is either empty or set to -1, no limit is applied. For a PATCH operation where a limit is configured, use a value of -1 to clear the limit.

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

users

Name	Type	Description
id	string	Quota target user ID
name	string	Quota target user name

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

quota_rule

Name	Type	Description
files	files	
space	space	
user_mapping	boolean	This parameter enables user mapping for user quota policy rules. This is valid in POST or PATCH for user quota policy rules only.
uuid	string	Unique identifier for the quota policy rule. This field is generated when the quota policy rule is created.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage shelves

Storage shelves endpoint overview

Retrieving storage shelf information

The storage shelf GET API retrieves all of the shelves in the cluster.

Examples

1) Retrieve a list of shelves from the cluster

The following example shows the response with a list of shelves in the cluster:

```
# The API:
/api/storage/shelves

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/shelves" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uid": "3109174803597886800",
      "_links": {
        "self": {
          "href": "/api/storage/shelves/3109174803597886800"
        }
      }
    }
  ],
}
```

```

{
  "uid": "9237728366621690448",
  "_links": {
    "self": {
      "href": "/api/storage/shelves/9237728366621690448"
    }
  }
},
{
  "uid": "9946762738829886800",
  "_links": {
    "self": {
      "href": "/api/storage/shelves/9946762738829886800"
    }
  }
},
{
  "uid": "10318311901725526608",
  "_links": {
    "self": {
      "href": "/api/storage/shelves/10318311901725526608"
    }
  }
},
{
  "uid": "13477584846688355664",
  "_links": {
    "self": {
      "href": "/api/storage/shelves/13477584846688355664"
    }
  }
}
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/shelves/"
  }
}
}

```

2) Retrieve a specific shelf from the cluster

The following example shows the response of the requested shelf. If there is no shelf with the requested uid, an error is returned.

```
# The API:
/api/storage/shelves/{uid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/shelves/3109174803597886800" -H
"accept: application/hal+json"

# The response:
{
  "uid": "3109174803597886800",
  "name": "6.10",
  "id": "10",
  "serial_number": "SHU0954292N0HAH",
  "model": "DS4246",
  "module_type": "iom6",
  "internal": false,
  "local": true,
  "state": "ok",
  "connection_type": "sas",
  "disk_count": 24,
  "paths": [
    {
      "name": "0e",
      "node": {
        "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
        "name": "node-1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0e"
        }
      }
    },
    {
      "name": "0g",
```

```

    "node": {
      "uuid": "0530d6c1-8c6d-11e8-907f-00a0985a72ee",
      "name": "node-1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0530d6c1-8c6d-11e8-907f-00a0985a72ee"
        }
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/ports/0530d6c1-8c6d-11e8-907f-00a0985a72ee/0g"
      }
    }
  ],
  "bays": [
    {
      "id": 0,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 1,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 2,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 3,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 4,
      "has_disk": true,

```

```
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 5,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 6,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 7,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 8,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 9,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 10,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
    "id": 11,
    "has_disk": true,
    "type": "single_disk",
    "state": "ok"
  },
  {
```

```
    "id": 12,  
    "has_disk": true,  
    "type": "single_disk",  
    "state": "ok"  
  },  
  {  
    "id": 13,  
    "has_disk": true,  
    "type": "single_disk",  
    "state": "ok"  
  },  
  {  
    "id": 14,  
    "has_disk": true,  
    "type": "single_disk",  
    "state": "ok"  
  },  
  {  
    "id": 15,  
    "has_disk": true,  
    "type": "single_disk",  
    "state": "ok"  
  },  
  {  
    "id": 16,  
    "has_disk": true,  
    "type": "single_disk",  
    "state": "ok"  
  },  
  {  
    "id": 17,  
    "has_disk": true,  
    "type": "single_disk",  
    "state": "ok"  
  },  
  {  
    "id": 18,  
    "has_disk": true,  
    "type": "single_disk",  
    "state": "ok"  
  },  
  {  
    "id": 19,  
    "has_disk": true,  
    "type": "single_disk",  
    "state": "ok"
```

```

    },
    {
      "id": 20,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 21,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 22,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    },
    {
      "id": 23,
      "has_disk": true,
      "type": "single_disk",
      "state": "ok"
    }
  ],
  "frus": [
    {
      "type": "module",
      "id": 0,
      "state": "ok",
      "part_number": "111-00690+B2",
      "serial_number": "8001900099",
      "firmware_version": "0191"
    },
    {
      "type": "module",
      "id": 1,
      "state": "ok",
      "part_number": "111-00190+B0",
      "serial_number": "7903785183",
      "firmware_version": "0191"
    },
    {
      "type": "psu",
      "id": 1,

```



```

    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW82562007513E",
    "firmware_version": "0311"
  },
  {
    "type": "psu",
    "id": 2,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW825620075138",
    "firmware_version": "0311"
  },
  {
    "type": "psu",
    "id": 3,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW8256200750BA",
    "firmware_version": "0311"
  },
  {
    "type": "psu",
    "id": 4,
    "state": "ok",
    "part_number": "0082562-12",
    "serial_number": "PMW8256200750A2",
    "firmware_version": "0311"
  }
],
"ports": [
  {
    "id": 0,
    "module_id": "a",
    "designator": "square",
    "state": "connected",
    "internal": false,
    "wwn": "500A098000C9EDBF",
    "cable": {
      "identifier": "5001086000702488-500a098000c9edbf",
      "part_number": "112-00430+A0",
      "length": "2m",
      "serial_number": "APF16510229807"
    },
    "remote": {
      "wwn": "5001086000702488",

```

```

    "phy": "08"
  }
},
{
  "id": 1,
  "module_id": "a",
  "designator": "circle",
  "state": "connected",
  "internal": false,
  "wnn": "500A098000C9EDBF",
  "cable": {
    "identifier": "500a098000d5c4bf-500a098000c9edbf",
    "part_number": "112-00176+A0",
    "length": "0.5-1.0m",
    "serial_number": "APF133917610YT"
  },
  "remote": {
    "wnn": "500A098000D5C4BF",
    "phy": "00"
  }
},
{
  "id": 2,
  "module_id": "b",
  "designator": "square",
  "state": "connected",
  "internal": false,
  "wnn": "500A098004F208BF",
  "cable": {
    "identifier": "5001086000702648-500a098004f208bf",
    "part_number": "112-00430+A0",
    "length": "2m",
    "serial_number": "APF16510229540"
  },
  "remote": {
    "wnn": "5001086000702648",
    "phy": "08"
  }
},
{
  "id": 3,
  "module_id": "b",
  "designator": "circle",
  "state": "connected",
  "internal": false,
  "wnn": "500A098004F208BF",

```

```

    "cable": {
      "identifier": "500a0980062ba33f-500a098004f208bf",
      "part_number": "112-00176+20",
      "length": "0.5-1.0m",
      "serial_number": "832210017"
    },
    "remote": {
      "wnn": "500A0980062BA33F",
      "phy": "00"
    }
  },
  "_links": {
    "self": {
      "href": "/api/storage/shelves/3109174803597886800"
    }
  }
}

```

Retrieve shelves

GET /storage/shelves

Introduced In: 9.6

Retrieves a collection of shelves.

Related ONTAP commands

- `storage shelf show`
- `storage shelf port show`
- `storage shelf drawer show`

Learn more

- [DOC /storage/shelves](#)

Parameters

Name	Type	In	Required	Description
paths.node.name	string	query	False	Filter by paths.node.name
paths.node.uuid	string	query	False	Filter by paths.node.uuid

Name	Type	In	Required	Description
paths.name	string	query	False	Filter by paths.name
connection_type	string	query	False	Filter by connection_type
frus.part_number	string	query	False	Filter by frus.part_number
frus.serial_number	string	query	False	Filter by frus.serial_number
frus.type	string	query	False	Filter by frus.type
frus.state	string	query	False	Filter by frus.state
frus.id	integer	query	False	Filter by frus.id
frus.firmware_version	string	query	False	Filter by frus.firmware_version
module_type	string	query	False	Filter by module_type
disk_count	integer	query	False	Filter by disk_count
bays.type	string	query	False	Filter by bays.type
bays.state	string	query	False	Filter by bays.state
bays.id	integer	query	False	Filter by bays.id
bays.has_disk	boolean	query	False	Filter by bays.has_disk
serial_number	string	query	False	Filter by serial_number
id	string	query	False	Filter by id
state	string	query	False	Filter by state

Name	Type	In	Required	Description
local	boolean	query	False	Filter by local <ul style="list-style-type: none"> • Introduced in: 9.8
drawers.part_number	string	query	False	Filter by drawers.part_number
drawers.error	string	query	False	Filter by drawers.error
drawers.serial_number	string	query	False	Filter by drawers.serial_number
drawers.disk_count	integer	query	False	Filter by drawers.disk_count
drawers.closed	boolean	query	False	Filter by drawers.closed
drawers.state	string	query	False	Filter by drawers.state
drawers.id	integer	query	False	Filter by drawers.id
vendor.serial_number	string	query	False	Filter by vendor.serial_number <ul style="list-style-type: none"> • Introduced in: 9.8
vendor.product	string	query	False	Filter by vendor.product <ul style="list-style-type: none"> • Introduced in: 9.8
vendor.manufacturer	string	query	False	Filter by vendor.manufacturer <ul style="list-style-type: none"> • Introduced in: 9.8

Name	Type	In	Required	Description
vendor.part_number	string	query	False	Filter by vendor.part_number • Introduced in: 9.8
ports.internal	boolean	query	False	Filter by ports.internal
ports.id	integer	query	False	Filter by ports.id
ports.state	string	query	False	Filter by ports.state
ports.module_id	string	query	False	Filter by ports.module_id
ports.remote.device	string	query	False	Filter by ports.remote.device • Introduced in: 9.8
ports.remote.phy	string	query	False	Filter by ports.remote.phy
ports.remote.wwn	string	query	False	Filter by ports.remote.wwn
ports.remote.chassis	string	query	False	Filter by ports.remote.chassis
ports.remote.mac_address	string	query	False	Filter by ports.remote.mac_address
ports.remote.port	string	query	False	Filter by ports.remote.port
ports.mac_address	string	query	False	Filter by ports.mac_address
ports.designator	string	query	False	Filter by ports.designator

Name	Type	In	Required	Description
ports.cable.part_number	string	query	False	Filter by ports.cable.part_number
ports.cable.identifier	string	query	False	Filter by ports.cable.identifier
ports.cable.length	string	query	False	Filter by ports.cable.length
ports.cable.serial_number	string	query	False	Filter by ports.cable.serial_number
ports.wwn	string	query	False	Filter by ports.wwn
internal	boolean	query	False	Filter by internal
model	string	query	False	Filter by model
name	string	query	False	Filter by name
uid	string	query	False	Filter by uid
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[shelf]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "bays": [
        {
          "id": "0",
          "state": "ok",
          "type": "single_disk"
        }
      ],
      "connection_type": "sas",
      "disk_count": "12",
      "drawers": [
        {
          "disk_count": "12",
          "error": "string",
          "part_number": "111-03071",
          "serial_number": "021604008263",
          "state": "ok"
        }
      ],
      "frus": [
        {
          "firmware_version": "0191",
          "part_number": "111-00690+A2",
          "serial_number": "8000166294",
          "state": "error",
          "type": "module"
        }
      ],
      "id": "1",
      "model": "DS2246",
      "module_type": "iom6",
      "name": "1.1",
      "paths": [
        {
```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "2a",
    "node": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "node1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "ports": [
    {
      "cable": {
        "identifier": "500a0980000b6c3f-50000d1703544b80",
        "length": "2m",
        "part_number": "112-00431+A0",
        "serial_number": "616930439"
      },
      "designator": "square",
      "id": "0",
      "mac_address": "string",
      "module_id": "a",
      "remote": {
        "chassis": "string",
        "device": "string",
        "mac_address": "string",
        "phy": "12",
        "port": "string",
        "wwn": "50000D1703544B80"
      },
      "state": "connected",
      "wwn": "500A0980000B6C3F"
    }
  ],
  "serial_number": "SHFMS1514000895",
  "state": "ok",
  "uid": "7777841915827391056",
  "vendor": {
    "manufacturer": "XYZ",

```

```
    "part_number": "A92831142733",
    "product": "LS2246",
    "serial_number": "891234572210221"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

bays

Name	Type	Description
has_disk	boolean	
id	integer	
state	string	
type	string	

drawers

Name	Type	Description
closed	boolean	
disk_count	integer	
error	string	
id	integer	
part_number	string	
serial_number	string	
state	string	

frus

Name	Type	Description
firmware_version	string	
id	integer	
part_number	string	
serial_number	string	
state	string	

Name	Type	Description
type	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

paths

Storage port

Name	Type	Description
_links	_links	
name	string	
node	node	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

remote

Name	Type	Description
chassis	string	
device	string	
mac_address	string	
phy	string	
port	string	
wwn	string	

ports

Name	Type	Description
cable	cable	
designator	string	
id	integer	
internal	boolean	
mac_address	string	
module_id	string	
remote	remote	
state	string	
wwn	string	

vendor

Name	Type	Description
manufacturer	string	Manufacturer name
part_number	string	Part number
product	string	Product name
serial_number	string	Serial number

shelf

Name	Type	Description
bays	array[bays]	
connection_type	string	
disk_count	integer	
drawers	array[drawers]	
frus	array[frus]	
id	string	
internal	boolean	
local	boolean	
model	string	
module_type	string	
name	string	

Name	Type	Description
paths	array[paths]	
ports	array[ports]	
serial_number	string	
state	string	
uid	string	
vendor	vendor	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a shelf

GET /storage/shelves/{uid}

Introduced In: 9.6

Retrieves a specific shelf.

Related ONTAP commands

- storage shelf show
- storage shelf port show
- storage shelf drawer show

Learn more

- [DOC /storage/shelves](#)

Parameters

Name	Type	In	Required	Description
uid	string	path	True	Shelf UID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
bays	array[bays]	
connection_type	string	
disk_count	integer	
drawers	array[drawers]	
frus	array[frus]	
id	string	
internal	boolean	
local	boolean	
model	string	
module_type	string	
name	string	
paths	array[paths]	
ports	array[ports]	
serial_number	string	
state	string	
uid	string	
vendor	vendor	

Example response

```
{
  "bays": [
    {
      "id": "0",
      "state": "ok",
      "type": "single_disk"
    }
  ],
  "connection_type": "sas",
  "disk_count": "12",
  "drawers": [
    {
      "disk_count": "12",
      "error": "string",
      "part_number": "111-03071",
      "serial_number": "021604008263",
      "state": "ok"
    }
  ],
  "frus": [
    {
      "firmware_version": "0191",
      "part_number": "111-00690+A2",
      "serial_number": "8000166294",
      "state": "error",
      "type": "module"
    }
  ],
  "id": "1",
  "model": "DS2246",
  "module_type": "iom6",
  "name": "1.1",
  "paths": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "2a",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        }
      }
    }
  ]
}
```

```

    }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"ports": [
  {
    "cable": {
      "identifier": "500a0980000b6c3f-50000d1703544b80",
      "length": "2m",
      "part_number": "112-00431+A0",
      "serial_number": "616930439"
    },
    "designator": "square",
    "id": "0",
    "mac_address": "string",
    "module_id": "a",
    "remote": {
      "chassis": "string",
      "device": "string",
      "mac_address": "string",
      "phy": "12",
      "port": "string",
      "wwn": "50000D1703544B80"
    },
    "state": "connected",
    "wwn": "500A0980000B6C3F"
  }
],
"serial_number": "SHFMS1514000895",
"state": "ok",
"uid": "7777841915827391056",
"vendor": {
  "manufacturer": "XYZ",
  "part_number": "A92831142733",
  "product": "LS2246",
  "serial_number": "891234572210221"
}
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

bays

Name	Type	Description
has_disk	boolean	
id	integer	
state	string	
type	string	

drawers

Name	Type	Description
closed	boolean	
disk_count	integer	
error	string	
id	integer	
part_number	string	
serial_number	string	
state	string	

frus

Name	Type	Description
firmware_version	string	
id	integer	
part_number	string	
serial_number	string	
state	string	
type	string	

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

paths

Storage port

Name	Type	Description
_links	_links	
name	string	
node	node	

cable

Name	Type	Description
identifier	string	
length	string	
part_number	string	
serial_number	string	

remote

Name	Type	Description
chassis	string	
device	string	
mac_address	string	
phy	string	
port	string	
wwn	string	

ports

Name	Type	Description
cable	cable	
designator	string	
id	integer	
internal	boolean	

Name	Type	Description
mac_address	string	
module_id	string	
remote	remote	
state	string	
wwn	string	

vendor

Name	Type	Description
manufacturer	string	Manufacturer name
part_number	string	Part number
product	string	Product name
serial_number	string	Serial number

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage Snapshot copy policies

Storage snapshot-policies endpoint overview

Overview

In ONTAP, scheduled Snapshot copy creation works based on Snapshot copy policies.

ONTAP provides three cluster-wide Snapshot copy policies: "default", "default-1weekly" and "none".

A Snapshot copy policy can have more than one schedule associated with it.

A Snapshot copy policy can be linked to a storage object and based on the schedule in the policy, Snapshot copies will be created on the object at that interval.

Each schedule in a Snapshot copy policy has a Snapshot copy name prefix attached to it. Every Snapshot copy created using this policy will have this prefix in its name.

There is also a retention count associated with every schedule. This count indicates the maximum number of Snapshot copies that can exist for a given schedule. Once the Snapshot copy count reaches the retention count, on the next create operation, the oldest Snapshot copy is deleted.

Snapshot copy policy APIs

The following APIs are used to perform operations related to Snapshot copy policy information:

– POST /api/storage/snapshot-policies

– GET /api/storage/snapshot-policies

– GET /api/storage/snapshot-policies/{uuid}

– PATCH /api/storage/snapshot-policies/{uuid}

– DELETE /api/storage/snapshot-policies/{uuid}

Examples

Creating a Snapshot copy policy

The POST operation is used to create a Snapshot copy policy with the specified attributes.

```
# The API:
/api/storage/snapshot-policies

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/snapshot-policies" -H
"accept: application/hal+json" -d '{"name": "new_policy", "enabled":
"true", "comment": "policy comment", "copies": [{ "schedule": { "name":
"5min" }, "count": "5", "prefix": "xyz" }], "svm": { "name": "vs0"}}'

# The response:
HTTP/1.1 201 Created
Date: Tue, 12 Mar 2019 21:20:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/snapshot-policies/a69d8173-450c-11e9-aa44-
005056bbc848
Content-Length: 369
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "uuid": "a69d8173-450c-11e9-aa44-005056bbc848",
      "svm": {
        "name": "vs0"
      },
      "name": "new_policy",
      "comment": "This is a 5min schedule policy",
      "enabled": true,
      "copies": [
        {
          "count": 5,
          "schedule": {
            "name": "5min"
          }
        }
      ]
    }
  ]
}
```

Retrieving Snapshot copy policy attributes

The GET operation is used to retrieve Snapshot copy policy attributes.


```
# The API:
/api/storage/snapshot-policies

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/" -H "accept:
application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:17:17 GMT
Server: libzapid-http
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 686
Content-Type: application/json
{
  "records": [
    {
      "uuid": "0fa7a554-348d-11e9-b55e-005056bbf1c8",
      "name": "spsv0",
      "_links": {
        "self": {
          "href": "/api/storage/snapshot-policies/0fa7a554-348d-11e9-b55e-
005056bbf1c8"
        }
      }
    },
    {
      "uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
      "name": "default",
      "_links": {
        "self": {
          "href": "/api/storage/snapshot-policies/3c112527-2fe8-11e9-b55e-
005056bbf1c8"
        }
      }
    },
    {
      "uuid": "3c1c1656-2fe8-11e9-b55e-005056bbf1c8",
      "name": "default-1weekly",
      "_links": {
        "self": {
          "href": "/api/storage/snapshot-policies/3c1c1656-2fe8-11e9-b55e-
005056bbf1c8"
        }
      }
    }
  ]
}
```

```

    }
  },
  {
    "uuid": "3c228b82-2fe8-11e9-b55e-005056bbf1c8",
    "name": "none",
    "_links": {
      "self": {
        "href": "/api/storage/snapshot-policies/3c228b82-2fe8-11e9-b55e-005056bbf1c8"
      }
    }
  }
],
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/storage/snapshot-policies/"
  }
}
}

```

Retrieving the attributes of a specific Snapshot copy policy

The GET operation is used to retrieve the attributes of a specific Snapshot copy policy.

```

# The API:
/api/storage/snapshot-policies

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/3c112527-2fe8-11e9-b55e-005056bbf1c8" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:24:48 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 381
Content-Type: application/json
{
  "uuid": "3c112527-2fe8-11e9-b55e-005056bbf1c8",
  "name": "default",
  "comment": "Default policy with hourly, daily & weekly schedules.",
  "enabled": true,
  "scope": "cluster",

```

```
"copies": [
  {
    "count": 6,
    "prefix": "hourly",
    "schedule": {
      "name": "hourly"
    }
  },
  {
    "count": 2,
    "prefix": "daily",
    "schedule": {
      "name": "daily"
    }
  },
  {
    "count": 2,
    "prefix": "weekly",
    "schedule": {
      "name": "weekly"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/storage/snapshot-policies/3c112527-2fe8-11e9-b55e-005056bbf1c8"
  }
}
```

Updating a Snapshot copy policy

The PATCH operation is used to update the specific attributes of a Snapshot copy policy.

```
# The API:
/api/storage/snapshot-policies/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/snapshot-policies/ae9e65c4-4506-11e9-aa44-005056bbc848" -d '{"enabled": "false" }' -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:27:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

Deleting a Snapshot copy policy

The DELETE operation is used to delete a Snapshot copy policy.

```
# The API:
/api/storage/snapshot-policies/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/snapshot-policies/ae9e65c4-4506-11e9-aa44-005056bbc848" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Tue, 12 Mar 2019 21:19:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

Retrieve Snapshot copy policies

GET /storage/snapshot-policies

Introduced In: 9.6

Retrieves a collection of Snapshot copy policies.

Related ONTAP commands

- `snapshot policy show`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
enabled	boolean	query	False	Filter by enabled
uuid	string	query	False	Filter by uuid
copies.snapmirror_label	string	query	False	Filter by copies.snapmirror_label
copies.schedule.name	string	query	False	Filter by copies.schedule.name
copies.count	integer	query	False	Filter by copies.count
copies.prefix	string	query	False	Filter by copies.prefix
name	string	query	False	Filter by name
comment	string	query	False	Filter by comment
scope	string	query	False	Filter by scope
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[snapshot_policy]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "copies": [
        {
          "prefix": "string",
          "schedule": {
            "name": "hourly"
          },
          "snapmirror_label": "string"
        }
      ],
      "enabled": 1,
      "name": "default",
      "scope": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
name	string	Schedule at which Snapshot copies are captured on the volume. Some common schedules already defined in the system are hourly, daily, weekly, at 15 minute intervals, and at 5 minute intervals. Snapshot copy policies with custom schedules can be referenced.

copies

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapshot_policy

The Snapshot copy policy object is associated with a read-write volume used to create and delete Snapshot copies at regular intervals.

Name	Type	Description
_links	_links	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Create a Snapshot copy policy

POST /storage/snapshot-policies

Introduced In: 9.6

Creates a Snapshot copy policy.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the Snapshot copy policy.
- `name` - Name for the Snapshot copy policy.
- `copies.schedule` - Schedule at which Snapshot copies are captured on the volume.
- `copies.count` - Number of Snapshot copies to maintain for this schedule.

Recommended optional properties

- `copies.prefix` - Prefix to use when creating Snapshot copies at regular intervals.

Default property values

If not specified in POST, the following default property values are assigned:

- `enabled` - *true*
- `copies.prefix` - Value of `schedule.name`

Related ONTAP commands

- `snapshot policy create`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	
uuid	string	

Example request

```
{
  "comment": "string",
  "copies": [
    {
      "prefix": "string",
      "schedule": {
        "name": "hourly"
      },
      "snapmirror_label": "string"
    }
  ],
  "enabled": 1,
  "name": "default",
  "scope": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1638407	When adding schedule to a Snapshot copy policy, the count for that schedule must be specified.
1638408	When adding schedule to a Snapshot copy policy, the schedule name must be specified.
1638413	Schedule not found.
1638417	Specified policy name is invalid.

Error Code	Description
1638451	This operation would result in total Snapshot copy count for the policy to exceed maximum supported count.
1638508	Another schedule has the same prefix within this policy.
1638526	This operation is not supported on a node Vserver.
1638527	Policy name already exists.
1638528	This operation is not supported in a mixed-version cluster.
1638531	This operation is not supported because specified policy is owned by the cluster admin.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

schedule

Name	Type	Description
name	string	Schedule at which Snapshot copies are captured on the volume. Some common schedules already defined in the system are hourly, daily, weekly, at 15 minute intervals, and at 5 minute intervals. Snapshot copy policies with custom schedules can be referenced.

copies

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapshot_policy

The Snapshot copy policy object is associated with a read-write volume used to create and delete Snapshot copies at regular intervals.

Name	Type	Description
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a Snapshot copy policy

DELETE /storage/snapshot-policies/{uuid}

Introduced In: 9.6

Deletes a Snapshot copy policy

Related ONTAP commands

- `snapshot policy delete`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Snapshot copy policy UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
1638415	Cannot delete policy. Reason: Policy is in use by at least one volume.
1638416	Cannot delete policy. Reason: Cannot verify whether policy is in use.
1638430	Cannot delete policy. Reason: Policy is in use by at least one Vserver.
1638430	Cannot delete built-in policy.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve Snapshot copy policy details

GET /storage/snapshot-policies/{uuid}

Introduced In: 9.6

Retrieves details of a specific Snapshot copy policy.

Related ONTAP commands

- `snapshot policy show`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Snapshot copy policy UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
comment	string	A comment associated with the Snapshot copy policy.
copies	array[copies]	
enabled	boolean	Is the Snapshot copy policy enabled?
name	string	Name of the Snapshot copy policy.
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
svm	svm	
uuid	string	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "copies": [
    {
      "prefix": "string",
      "schedule": {
        "name": "hourly"
      },
      "snapmirror_label": "string"
    }
  ],
  "enabled": 1,
  "name": "default",
  "scope": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
name	string	Schedule at which Snapshot copies are captured on the volume. Some common schedules already defined in the system are hourly, daily, weekly, at 15 minute intervals, and at 5 minute intervals. Snapshot copy policies with custom schedules can be referenced.

copies

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.

Name	Type	Description
uuid	string	The unique identifier of the SVM.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a Snapshot copy policy

PATCH /storage/snapshot-policies/{uuid}

Introduced In: 9.6

Updates a Snapshot copy policy

Related ONTAP commands

- `snapshot policy modify`
- `snapshot policy modify-schedule`
- `snapshot policy add-schedule`

Learn more

- [DOC /storage/snapshot-policies](#)

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Snapshot copy policy UUID

Request Body

Name	Type	Description
comment	string	A comment associated with the Snapshot copy policy.
enabled	boolean	Is the Snapshot copy policy enabled?
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
uuid	string	

Example request

```
{
  "comment": "string",
  "enabled": 1,
  "scope": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
1638414	Cannot enable policy. Reason: Specified schedule not found.

See Definitions

href

Name	Type	Description
href	string	

_links

schedule

Name	Type	Description
name	string	Schedule at which Snapshot copies are captured on the volume. Some common schedules already defined in the system are hourly, daily, weekly, at 15 minute intervals, and at 5 minute intervals. Snapshot copy policies with custom schedules can be referenced.

copies

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snapshot_policy

The Snapshot copy policy object is associated with a read-write volume used to create and delete Snapshot copies at regular intervals.

Name	Type	Description
comment	string	A comment associated with the Snapshot copy policy.
enabled	boolean	Is the Snapshot copy policy enabled?
scope	string	Set to "svm" when the request is on a data SVM, otherwise set to "cluster".
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage Snapshot copy policies and schedules

Storage snapshot-policies snapshot_policy.uuid schedules endpoint overview

Overview

In ONTAP, scheduled Snapshot copy creation works based on the schedules associated with Snapshot copy policies.

ONTAP provides six cluster-wide schedules: "5min", "8hour", "hourly", "daily", "weekly" and "monthly".

A Snapshot copy policy is created using at least one of these schedules and up to 5 schedules can be associated with a Snapshot copy policy.

A Snapshot copy policy can be linked to a storage object and based on the schedule in the policy, Snapshot copies are created on the object at that interval.

Each schedule in a Snapshot copy policy has a Snapshot copy name prefix attached to it. Every Snapshot

copy created using this policy has this prefix in its name.

There is also a retention count associated with every schedule. This count indicates the maximum number of Snapshot copies that can exist for a given schedule.

Once the Snapshot copy count reaches the retention count, on the next create operation, the oldest Snapshot copy is deleted.

A schedule can be added, modified or deleted from a Snapshot copy policy.

Snapshot copy policy schedule APIs

The following APIs are used to perform operations related to Snapshot copy policy schedules:

– POST /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

– GET /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

– GET /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

– PATCH /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

– DELETE /api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

Examples

Adding schedule to a Snapshot copy policy

The POST operation is used to create a schedule for a Snapshot copy policy with the specified attributes.

```
# The API:
/api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-818e-11e9-b4f4-005056bbab9c/schedules" -H "accept: application/hal+json" -d '{"schedule.uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c", "count": "5", "prefix": "new_hourly" }'

# The response:
HTTP/1.1 201 Created
Date: Wed, 29 May 2019 22:41:33 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/snapshot-policies/32a0841a-818e-11e9-b4f4-005056bbab9c/schedules
Content-Length: 271
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c"
      },
      "count": 5,
      "prefix": "new_monthly"
    }
  ]
}
```

Retrieving Snapshot copy policy schedules

The GET operation is used to retrieve Snapshot copy policy schedules.

```
# The API:
/api/storage/snapshot-policies/{snapshot_policy.uuid}/schedules/

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-818e-11e9-b4f4-005056bbab9c/schedules" -H "accept: application/hal+json"
```

```
# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 22:49:58 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 898
Content-Type: application/json
{
  "records": [
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "63d017dc-818a-11e9-b4f4-005056bbab9c",
        "name": "5min"
      }
    },
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "64a5c5da-818a-11e9-b4f4-005056bbab9c",
        "name": "8hour"
      }
    },
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "63e21a3e-818a-11e9-b4f4-005056bbab9c",
        "name": "daily"
      }
    },
    {
      "snapshot_policy": {
        "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"
      },
      "schedule": {
        "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c",
        "name": "monthly"
      }
    }
  ]
}
```

```
}  
],  
"num_records": 4  
}
```

Retrieving the attributes of a specific Snapshot copy policy schedule

The GET operation is used to retrieve the attributes of a specific Snapshot copy policy schedule.

```
# The API:  
/api/storage/snapshot-  
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}  
  
# The call:  
curl -X GET "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-  
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-  
005056bbab9c" -H "accept: application/hal+json"  
  
# The response:  
HTTP/1.1 200 OK  
Date: Wed, 29 May 2019 22:54:06 GMT  
Server: libzapid-httpd  
X-Content-Type-Options: nosniff  
Cache-Control: no-cache,no-store,must-revalidate  
Content-Length: 238  
Content-Type: application/json  
{  
  "snapshot_policy": {  
    "uuid": "32a0841a-818e-11e9-b4f4-005056bbab9c"  
  },  
  "schedule": {  
    "uuid": "7c985d80-818a-11e9-b4f4-005056bbab9c",  
    "name": "monthly"  
  },  
  "count": 5,  
  "prefix": "new_monthly",  
  "snapmirror_label": "--"  
}
```

Updating a Snapshot copy policy schedule

The PATCH operation is used to update the specific attributes of a Snapshot copy policy.

```
# The API:
/api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-
005056bbab9c" -d '{"count": "10" }' -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 23:08:00 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

Deleting a Snapshot copy policy

The DELETE operation is used to delete a Snapshot copy policy.

```
# The API:
/api/storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/snapshot-policies/32a0841a-
818e-11e9-b4f4-005056bbab9c/schedules/7c985d80-818a-11e9-b4f4-
005056bbab9c" -H "accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 29 May 2019 23:12:32 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/json
```

Retrieve Snapshot copy policy schedules

GET /storage/snapshot-policies/{snapshot_policy.uuid}/schedules

Introduced In: 9.8

Retrieves a collection of Snapshot copy policy schedules.

Related ONTAP commands

- `snapshot policy show`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
snapshot_policy.uuid	string	path	True	Snapshot copy policy UUID
prefix	string	query	False	Filter by prefix
count	integer	query	False	Filter by count
schedule.name	string	query	False	Filter by schedule.name
schedule.uuid	string	query	False	Filter by schedule.uuid
snapshot_policy.name	string	query	False	Filter by snapshot_policy.name
snapshot_policy.uuid	string	query	False	Filter by snapshot_policy.uuid
snapmirror_label	string	query	False	Filter by snapmirror_label
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[snapshot_policy_schedule]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "prefix": "string",
      "schedule": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "weekly",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "snapmirror_label": "string",
      "snapshot_policy": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "default",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

snapshot_policy_schedule

The Snapshot copy policy schedule object is associated with a Snapshot copy policy and it defines the interval at which Snapshot copies are created and deleted.

Name	Type	Description
_links	_links	

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Add a schedule to a Snapshot copy policy

POST /storage/snapshot-policies/{snapshot_policy.uuid}/schedules

Introduced In: 9.8

Adds a schedule to a Snapshot copy policy.

Required properties

- `schedule.uuid` or `schedule.name` - Schedule at which Snapshot copies are captured on the volume.

- `count` - Number of Snapshot copies to maintain for this schedule.

Recommended optional properties

- `prefix` - Prefix to use when creating Snapshot copies at regular intervals.

Default property values

If not specified in POST, the following default property values are assigned:

- `prefix` - Value of `schedule.name`

Related ONTAP commands

- `snapshot policy add-schedule`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
<code>snapshot_policy.uuid</code>	string	path	True	Snapshot copy policy UUID
<code>return_records</code>	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
<code>count</code>	integer	The number of Snapshot copies to maintain for this schedule.
<code>prefix</code>	string	The prefix to use while creating Snapshot copies at regular intervals.
<code>schedule</code>	schedule	
<code>snapmirror_label</code>	string	Label for SnapMirror operations
<code>snapshot_policy</code>	snapshot_policy	This is a reference to the Snapshot copy policy.

Example request

```
{
  "prefix": "string",
  "schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapmirror_label": "string",
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1638407	When adding schedule to a Snapshot copy policy, the count for that schedule must be specified.
1638410	Specified schedule already exists in snapshot policy.
1638413	Schedule not found.
1638451	This operation would result in total Snapshot copy count for the policy to exceed maximum supported count.
1638508	Another schedule has the same prefix within this policy.
1638528	This operation is not supported in a mixed-version cluster.
1638531	This operation is not supported because specified policy is owned by the cluster admin.

See Definitions

href

Name	Type	Description
href	string	

_links

schedule

Name	Type	Description
name	string	Job schedule name
uuid	string	Job schedule UUID

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
name	string	
uuid	string	

snapshot_policy_schedule

The Snapshot copy policy schedule object is associated with a Snapshot copy policy and it defines the interval at which Snapshot copies are created and deleted.

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a schedule from a Snapshot copy policy

```
DELETE /storage/snapshot-
policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}
```

Introduced In: 9.8

Deletes a schedule from a Snapshot copy policy

Related ONTAP commands

- `snapshot policy remove-schedule`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
snapshot_policy.uuid	string	path	True	Snapshot copy policy UUID
schedule.uuid	string	path	True	Snapshot copy policy schedule UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
1638412	Schedule does not exist in Snapshot policy.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve Snapshot copy policy schedule details

GET /storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

Introduced In: 9.8

Retrieves details of a specific Snapshot copy policy schedule.

Related ONTAP commands

- `snapshot policy show`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
snapshot_policy.uuid	string	path	True	Snapshot copy policy UUID
schedule.uuid	string	path	True	Snapshot copy policy schedule ID

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
count	integer	The number of Snapshot copies to maintain for this schedule.
prefix	string	The prefix to use while creating Snapshot copies at regular intervals.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "prefix": "string",
  "schedule": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapmirror_label": "string",
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

schedule

Name	Type	Description
_links	_links	
name	string	Job schedule name
uuid	string	Job schedule UUID

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code

Name	Type	Description
message	string	Error message
target	string	The target parameter that caused the error.

Update a Snapshot copy policy schedule

PATCH /storage/snapshot-policies/{snapshot_policy.uuid}/schedules/{schedule.uuid}

Introduced In: 9.8

Updates a Snapshot copy policy schedule

Related ONTAP commands

- `snapshot policy modify-schedule`

Learn more

- [DOC /storage/snapshot-policies/{snapshot_policy.uuid}/schedules](#)

Parameters

Name	Type	In	Required	Description
snapshot_policy.uuid	string	path	True	Snapshot copy policy UUID
schedule.uuid	string	path	True	Snapshot copy policy schedule UUID

Request Body

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

Example request

```
{
  "schedule": {
    "name": "weekly",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "snapmirror_label": "string",
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
1638451	This operation would result in total Snapshot copy count for the policy to exceed maximum supported count.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

schedule

Name	Type	Description
name	string	Job schedule name
uuid	string	Job schedule UUID

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
name	string	
uuid	string	

snapshot_policy_schedule

The Snapshot copy policy schedule object is associated with a Snapshot copy policy and it defines the interval at which Snapshot copies are created and deleted.

Name	Type	Description
count	integer	The number of Snapshot copies to maintain for this schedule.
schedule	schedule	
snapmirror_label	string	Label for SnapMirror operations
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a volume efficiency policy

POST /storage/volume-efficiency-policies

Introduced In: 9.8

Creates a volume efficiency policy.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the volume efficiency policy.
- `name` - Name for the volume efficiency policy.

Recommended optional properties

- `type` - Type of volume policy.
- `schedule` - Schedule the volume efficiency defined in minutes, hourly, daily and weekly.
- `duration` - Indicates the allowed duration for a session for policy type "threshold" and "auto".
- `start_threshold_percent` - Indicates the start threshold percentage for the policy type "threshold". It is mutually exclusive of the schedule.
- `qos_policy` - QoS policy for the sis operation.
- `comment` - A comment associated with the volume efficiency policy.
- `enabled` - Is the volume efficiency policy enabled?

Default property values

If not specified in POST, the following default property values are assigned:

- type - auto
- enabled - true
- qos_policy - best_effort

Related ONTAP commands

- volume efficiency policy create

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
comment	string	A comment associated with the volume efficiency policy.
duration	integer	This field is used with the policy types "scheduled" and "auto" to indicate the allowed duration for a session, in hours. Possible value is a number between 0 and 999 inclusive. Default is unlimited indicated by value 0.
enabled	boolean	Is the volume efficiency policy enabled?
name	string	Name of the volume efficiency policy.
qos_policy	string	QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations.

Name	Type	Description
schedule	schedule	
start_threshold_percent	integer	This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule.
svm	svm	SVM, applies only to SVM-scoped objects.
type	string	Type of volume efficiency policy.
uuid	string	Unique identifier of volume efficiency policy.

Example request

```
{
  "comment": "string",
  "duration": "5",
  "enabled": 1,
  "name": "default",
  "qos_policy": "string",
  "schedule": {
    "name": "daily"
  },
  "start_threshold_percent": "30",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "type": "string",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
}
```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
6881341	Specified schedule not found.
6881344	Failed to queue specified job.
6881345	This operation is not permitted on a node SVM.
6881349	Policy name is not valid.
6881362	Threshold percentage cannot be less than 1 percent.
6881433	For "{{0}}" type policy, attribute "{{1}}" is not supported.
6881435	Only a policy of type "threshold" can set the "start-threshold-percent" attribute.
6881436	For a policy of type "scheduled", a valid "schedule" is a required attribute.
6881454	An efficiency policy of type "threshold" requires an effective cluster version of ONTAP 8.3 or later.
6881474	Duration cannot be null.
6881475	Duration is not valid.
6881476	Duration cannot be less than 1 hour.
6881477	Duration cannot be more than 999 hours.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

schedule

Name	Type	Description
name	string	Schedule at which volume efficiency policies are captured on the SVM. Some common schedules already defined in the system are hourly, daily, weekly, at 5 minute intervals, and at 8 hour intervals. Volume efficiency policies with custom schedules can be referenced.

svm

SVM, applies only to SVM-scoped objects.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume_efficiency_policy

Name	Type	Description
comment	string	A comment associated with the volume efficiency policy.
duration	integer	This field is used with the policy types "scheduled" and "auto" to indicate the allowed duration for a session, in hours. Possible value is a number between 0 and 999 inclusive. Default is unlimited indicated by value 0.

Name	Type	Description
enabled	boolean	Is the volume efficiency policy enabled?
name	string	Name of the volume efficiency policy.
qos_policy	string	QoS policy for the sis operation. Possible values are background and best_effort. In background, sis operation will run in background with minimal or no impact on data serving client operations. In best_effort, sis operations may have some impact on data serving client operations.
schedule	schedule	
start_threshold_percent	integer	This field is used with the policy type "threshold" to indicate the threshold percentage for triggering the volume efficiency policy. It is mutually exclusive of the schedule.
svm	svm	SVM, applies only to SVM-scoped objects.
type	string	Type of volume efficiency policy.
uuid	string	Unique identifier of volume efficiency policy.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a volume efficiency policy

DELETE /storage/volume-efficiency-policies/{uuid}

Introduced In: 9.8

Deletes a volume efficiency policy.

Related ONTAP commands

- `volume efficiency policy modify`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Volume efficiency policy UUID

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Code

Error Code	Description
6881346	The policy was not deleted because the policy is in use by at least one volume.
6881347	This operation cannot be performed because the specified policy is owned by the cluster admin.

Error Code	Description
6881431	The specified policy is a predefined policy and cannot be deleted.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage storage volumes

Storage volumes endpoint overview

Overview

FlexVol volumes are logical containers used by ONTAP to serve data to clients. They contain file systems in a NAS environment and LUNs in a SAN environment.

A FlexGroup volume is a scale-out NAS container that provides high performance along with automatic load distribution and scalability. A FlexGroup volume contains several constituents that automatically and transparently share the traffic.

FlexClone volumes are writable, point-in-time copies of a FlexVol volume. At this time, FlexClones of FlexGroups are not supported.

Volumes with SnapLock type Compliance or Enterprise, are referred to as SnapLock volumes. Volumes with SnapLock type cannot be of FlexGroup style. Once a SnapLock aggregate is created, by default, volumes created inside the aggregate inherit the "snaplock" property from the aggregate. It is possible to create a SnapLock volume by specifying SnapLock parameters. SnapLock parameters are only available at the "advanced" privilege level.

ONTAP storage APIs allow you to create, modify, and monitor volumes and aggregates.

Storage efficiency

Storage efficiency is used to remove duplicate blocks in the data and to compress the data. Efficiency has

deduplication, compression, cross volume deduplication, and compaction options. On All Flash systems, all efficiencies are enabled by default on volume creation. Options such as "background/inline/both" are treated as both, which means both background and inline are enabled for any efficiency option. The option "none" disables both background and inline efficiency.

To enable any efficiency option on all-flash or FAS systems, background deduplication is always enabled.

Quotas

Quotas provide a way to restrict or track the files and space usage by a user, group, or qtree. Quotas are enabled for a specific FlexVol or a FlexGroup volume.

The following APIs can be used to enable or disable and obtain quota state for a FlexVol or a FlexGroup volume:

– PATCH /api/storage/volumes/{uuid} -d '{"quota.enabled":"true"}

– PATCH /api/storage/volumes/{uuid} -d '{"quota.enabled":"false"}

– GET /api/storage/volumes/{uuid}?fields=quota.state

File System Analytics

File system analytics provide a quick method for obtaining information summarizing properties of all files within any directory tree of a volume. For more information on file system analytics, see [DOC /storage/volumes{volume.uuid}/files/{path}](#) . Analytics can be enabled or disabled on individual volumes.

The following APIs can be used to enable or disable and obtain analytics state for a FlexVol volume or a FlexGroup volume:

– PATCH /api/storage/volumes/{uuid} -d '{"analytics.state":"on"}

– PATCH /api/storage/volumes/{uuid} -d '{"analytics.state":"off"}

– GET /api/storage/volumes/{uuid}?fields=analytics

QoS

QoS policy and settings enforce Service Level Objectives (SLO) on a volume. SLO can be set by specifying qos.max_throughput_iops and/or qos.max_throughput_mbps or qos.min_throughput_iops. Specifying min_throughput_iops is only supported on volumes hosted on a node that is flash optimized. A pre-created QoS policy can also be used by specifying qos.name or qos.uuid property.

Performance monitoring

Performance of a volume can be monitored by the `metric.*` and `statistics.*` fields. These show the performance of the volume in terms of IOPS, latency and throughput. The `metric.*` fields denote an average whereas `statistics.*` fields denote a real-time monotonically increasing value aggregated across all nodes.

Volume APIs

The following APIs are used to perform operations related with FlexVol volumes and FlexGroup volumes:

– POST /api/storage/volumes

– GET /api/storage/volumes

– GET /api/storage/volumes/{uuid}

– PATCH /api/storage/volumes/{uuid}

– DELETE /api/storage/volumes/{uuid}

Examples

Creating a volume

The POST request is used to create a new volume and to specify its properties.

```
# The API:
/api/storage/volumes

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name": "vol1",
"aggregates":[{"name":"aggr1"}], "svm":{"name" : "vs1"}}'

# The response:
{
  "job": {
    "uuid": "b89bc5dd-94a3-11e8-a7a3-0050568edf84",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/b89bc5dd-94a3-11e8-a7a3-0050568edf84"
      }
    }
  }
}
```

Creating a SnapLock volume and specifying its properties using POST

```
# The API:
/api/storage/volumes

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name": "vol1", "aggregates":[{"name":
"aggr1"}], "svm":{"name" : "vs1"}, "snaplock":{"retention":{"default":
"P20Y"}}}'

# The response:
{
  "job": {
    "uuid": "e45b123b-c228-11e8-aa20-0050568e36bb",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/e45b123b-c228-11e8-aa20-0050568e36bb"
      }
    }
  }
}
```

Creating a FlexGroup volume and specifying its properties using POST


```

# The API:
/api/storage/volumes

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name" : "voll", "state" : "online", "type" :
"RW", "aggregates" : [{"name" : "aggr1"}, {"name" : "aggr2"},
{"name":"aggr3"}], "constituents_per_aggregate" : "1", "svm" : {"name" :
"vs1"}, "size" : "240MB", "encryption" : {"enabled" : "False"},
"efficiency" : {"compression" : "both"}, "autosize" : {"maximum" :
"500MB", "minimum" : "240MB"}}'

# The response:
{
  "job": {
    "uuid": "3cfa38bd-3a78-11e9-ae39-0050568ed7dd",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/3cfa38bd-3a78-11e9-ae39-0050568ed7dd"
      }
    }
  }
}

```

Creating a FlexClone and specifying its properties using POST

```
# The API:
/api/storage/volumes

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json" -d '{"name":"voll_clone", "clone": {"parent_volume":
{"name": "voll"}, "is_flexclone": "true"},"svm":{"name": "vs0"}}'

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 26 Feb 2019 09:06:22 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volumes/?name=voll_clone
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "c9ee0040-39a5-11e9-9b24-00a098439a83",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/c9ee0040-39a5-11e9-9b24-00a098439a83"
      }
    }
  }
}
```

Volumes reported in the GET REST API

The following types of volumes are reported:

– RW, DP and LS volume

– FlexGroup volume

– FlexCache volume

– FlexClone volume

The following types of volumes are not reported:

– DEL volume

– TEMP volume

– Node Root volume

– System Vserver volume

– FlexGroup constituent

– FlexCache constituent

Examples

Retrieving the list of volumes

```
# The API:
/api/storage/volumes

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes" -H "accept:
application/hal+json"

# The response:
{
  "records": [
    {
      "uuid": "2d1167cc-c3f2-495a-a23f-8f50b071b9b8",
      "name": "vsdata_root",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/2d1167cc-c3f2-495a-a23f-
8f50b071b9b8"
        }
      }
    },
    {
      "uuid": "3969be7e-78b4-4b4c-82a4-fa86331f03df",
      "name": "vsfg_root",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/3969be7e-78b4-4b4c-82a4-
fa86331f03df"
        }
      }
    },
    {
      "uuid": "59c03ac5-e708-4ce8-a676-278dc249fda2",
      "name": "svm_root",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/59c03ac5-e708-4ce8-a676-
278dc249fda2"
        }
      }
    }
  ]
}
```

```

    }
  },
  {
    "uuid": "6802635b-8036-11e8-aae5-0050569503ac",
    "name": "fgvol",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/6802635b-8036-11e8-aae5-0050569503ac"
      }
    }
  },
  {
    "uuid": "d0c3359c-5448-4a9b-a077-e3295a7e9057",
    "name": "datavol",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/d0c3359c-5448-4a9b-a077-e3295a7e9057"
      }
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/volumes"
  }
}
}

```

Retrieving the attributes of a volume

The GET request is used to retrieve the attributes of a volume.

```

# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-a077-e3295a7e9057" -H "accept: application/hal+json"

# The response:
{
  "uuid": "d0c3359c-5448-4a9b-a077-e3295a7e9057",

```

```

"comment": "This is a data volume",
"create_time": "2018-07-05T14:56:44+05:30",
"language": "en_us",
"name": "datavol",
"size": 20971520,
"state": "online",
"style": "flexvol",
"tiering_policy": "auto",
"type": "rw",
"aggregates": [
  {
    "name": "data",
    "uuid": "aa742322-36bc-4d98-bbc4-0a827534c035",
    "_links": {
      "self": {
        "href": "/api/cluster/aggregates/data"
      }
    }
  }
],
"encryption": {
  "enabled": false,
  "state": "none",
  "key_id": "",
  "type": "none"
},
"error_state": {
  "has_bad_blocks": false,
  "is_inconsistent": false
},
"files": {
  "maximum": 566,
  "used": 96
},
"nas": {
  "gid": 2468,
  "security_style": "unix",
  "uid": 1357,
  "unix_permissions": 4755,
  "export_policy": {
    "name": "default",
    "id": 8589934593
  }
},
"metric": {
  "timestamp": "2019-04-09T05:50:15Z",

```

```

"status": "ok",
"duration": "PT15S",
"latency": {
  "other": 0,
  "total": 0,
  "read": 0,
  "write": 0
},
"iops": {
  "read": 0,
  "write": 0,
  "other": 0,
  "total": 0
},
"throughput": {
  "read": 0,
  "write": 0,
  "other": 0,
  "total": 0
},
"cloud": {
  "timestamp": "2019-04-09T05:50:15Z",
  "status": "ok",
  "duration": "PT15S",
  "iops" : {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "latency": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"flexcache": {
  "timestamp": "2019-04-09T05:50:15Z",
  "status": "ok",
  "duration": "PT1D",
  "cache_miss_percent": 0
},
"statistics": {
  "timestamp": "2019-04-09T05:50:42Z",

```

```

"status": "ok",
"latency_raw": {
  "other": 38298,
  "total": 38298,
  "read": 0,
  "write": 0
},
"iops_raw": {
  "read": 0,
  "write": 0,
  "other": 3,
  "total": 3
},
"throughput_raw": {
  "read": 0,
  "write": 0,
  "other": 0,
  "total": 0
},
"cloud": {
  "timestamp": "2019-04-09T05:50:42Z",
  "status": "ok",
  "iops_raw" : {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  },
  "latency_raw": {
    "read": 0,
    "write": 0,
    "other": 0,
    "total": 0
  }
},
"flexcache_raw": {
  "timestamp": "2019-04-09T05:50:15Z",
  "status": "ok",
  "cache_miss_blocks": 0,
  "client_requested_blocks": 0
},
"qos": {
  "policy": {
    "min_throughput_iops": 0,
    "max_throughput_iops": 1000,

```

```

    "max_throughput_mbps": 0,
    "uuid": "228454af-5a8b-11e9-bd5b-005056ac6f1f",
    "name": "pg1"
  }
},
"snaplock": {
  "append_mode_enabled": false,
  "autocommit_period": "none",
  "compliance_clock_time": "2019-05-24T10:59:00+05:30",
  "expiry_time": "2038-01-19T08:44:28+05:30",
  "is_audit_log": false,
  "litigation_count": 0,
  "privileged_delete": "disabled",
  "type": "enterprise",
  "retention": {
    "default": "P0Y",
    "minimum": "P0Y",
    "maximum": "P30Y"
  }
},
"snapshot_policy": {
  "name": "default"
},
"svm": {
  "name": "vsdata",
  "uuid": "d61b69f5-7458-11e8-ad3f-0050569503ac"
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/d0c3359c-5448-4a9b-a077-e3295a7e9057"
  }
}
}

```

Retrieving the quota state of a FlexVol or a FlexGroup volume


```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717/?fields=quota.state" -H "accept: application/hal+json"

# The response:
{
  "uuid": "cb20da45-4f6b-11e9-9a71-005056a7f717",
  "name": "fv",
  "quota": {
    "state": "on"
  },
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cb20da45-4f6b-11e9-9a71-005056a7f717/"
    }
  }
}
```

Updating the attributes of a volume

Examples

Updating the attributes of a volume

The PATCH request is used to update the attributes of a volume.

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{ "size": 26214400, "nas": {"security_style":
"mixed"}, "comment": "This is a data volume" }' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 31 Jul 2018 09:36:43 GMT
Server: libzapid-httpd
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "3c5be5a6-94a5-11e8-8ca3-00505695c11b",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/3c5be5a6-94a5-11e8-8ca3-00505695c11b"
      }
    }
  }
}
```

Updating the attributes of a FlexClone using PATCH

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"clone":{"split_initiated":"true"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "8e01747f-38e5-11e9-8a3a-00a09843994b",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/8e01747f-38e5-11e9-8a3a-00a09843994b"
      }
    }
  }
}
```

Enabling quotas for a FlexVol or a FlexGroup volume using PATCH

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"quota":{"enabled":"true"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "d2fe7299-57d0-11e9-a2dc-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/d2fe7299-57d0-11e9-a2dc-005056a7f717"
      }
    }
  }
}
```

Disabling quotas for a FlexVol or a FlexGroup volume using PATCH

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"quota":{"enabled":"false"}}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Mon, 25 Feb 2019 10:10:19 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "0c8f6bea-57d1-11e9-a2dc-005056a7f717",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/0c8f6bea-57d1-11e9-a2dc-005056a7f717"
      }
    }
  }
}
```

Add tiering object tags for a FlexVol using PATCH

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"tiering.object_tags": [ "key1=val1", "key2=val2"
]}' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Tue, 11 Feb 2020 19:29:25 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "d05012de-4d04-11ea-836b-005056bb6f9d",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/d05012de-4d04-11ea-836b-005056bb6f9d"
      }
    }
  }
}
```

Remove tiering object tags for a FlexVol using PATCH

```
# The API:
/api/storage/volumes/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/d0c3359c-5448-4a9b-
a077-e3295a7e9057" -d '{"tiering.object_tags": []}' -H "accept:
application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Fri, 24 Jan 2020 22:28:04 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "ca234df1-3ef8-11ea-9a56-005056bb69a1",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/ca234df1-3ef8-11ea-9a56-005056bb69a1"
      }
    }
  }
}
```

Deleting a volume

Example

Deleting a volume

The DELETE request is used to delete a volume.

```
# The API:
/api/storage/volumes

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/{uuid}" -H
"accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
cache-control: no-cache,no-store,must-revalidate
connection: Keep-Alive
content-length: 189
content-type: application/json
date: Wed, 01 Aug 2018 09:40:36 GMT
keep-alive: timeout=5, max=100
server: libzapid-httpd
{
  "job": {
    "uuid": "f1aa3eb8-956e-11e8-86bf-0050568e2249",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/f1aa3eb8-956e-11e8-86bf-0050568e2249"
      }
    }
  }
}
```

Retrieve volumes

GET /storage/volumes

Introduced In: 9.6

Retrieves volumes.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `is_svm_root`
- `analytics.*`
- `application.*`
- `encryption.*`

- clone.parent_snapshot.name
- clone.parent_snapshot.uuid
- clone.parent_svm.name
- clone.parent_svm.uuid
- clone.parent_volume.name
- clone.parent_volume.uuid
- clone.split_complete_percent
- clone.split_estimate
- clone.split_initiated
- efficiency.*
- error_state.*
- files.*
- nas.export_policy.id
- nas.gid
- nas.path
- nas.security_style
- nas.uid
- nas.unix_permissions
- snaplock.*
- restore_to.*
- snapshot_policy.uuid
- quota.*
- qos.*
- flexcache_endpoint_type
- space.block_storage_inactive_user_data
- space.capacity_tier_footprint
- space.performance_tier_footprint
- space.local_tier_footprint
- space.footprint
- space.over_provisioned
- space.metadata
- space.total_footprint
- space.logical_space.*
- space.snapshot.*

- `guarantee.*`
- `autosize.*`
- `movement.*`
- `statistics.*`

Related ONTAP commands

- `volume show`
- `volume clone show`
- `volume efficiency show`
- `volume encryption show`
- `volume flexcache show`
- `volume flexgroup show`
- `volume move show`
- `volume quota show`
- `volume show-space`
- `volume snaplock show`

Parameters

Name	Type	In	Required	Description
<code>is_object_store</code>	boolean	query	False	Filter by <code>is_object_store</code> <ul style="list-style-type: none"> • Introduced in: 9.8
<code>application.name</code>	string	query	False	Filter by <code>application.name</code>
<code>application.uuid</code>	string	query	False	Filter by <code>application.uuid</code>
<code>state</code>	string	query	False	Filter by <code>state</code>
<code>efficiency.compression</code>	string	query	False	Filter by <code>efficiency.compression</code>
<code>efficiency.cross_volume_dedupe</code>	string	query	False	Filter by <code>efficiency.cross_volume_dedupe</code>

Name	Type	In	Required	Description
efficiency.dedupe	string	query	False	Filter by efficiency.dedupe
efficiency.compaction	string	query	False	Filter by efficiency.compaction
efficiency.policy.name	string	query	False	Filter by efficiency.policy.name • Introduced in: 9.7
files.used	integer	query	False	Filter by files.used
files.maximum	integer	query	False	Filter by files.maximum
statistics.iops_raw.total	integer	query	False	Filter by statistics.iops_raw.total
statistics.iops_raw.read	integer	query	False	Filter by statistics.iops_raw.read
statistics.iops_raw.other	integer	query	False	Filter by statistics.iops_raw.other
statistics.iops_raw.write	integer	query	False	Filter by statistics.iops_raw.write
statistics.timestamp	string	query	False	Filter by statistics.timestamp
statistics.latency_raw.total	integer	query	False	Filter by statistics.latency_raw.total
statistics.latency_raw.read	integer	query	False	Filter by statistics.latency_raw.read

Name	Type	In	Required	Description
statistics.latency_raw.other	integer	query	False	Filter by statistics.latency_raw.other
statistics.latency_raw.write	integer	query	False	Filter by statistics.latency_raw.write
statistics.status	string	query	False	Filter by statistics.status
statistics.cloud.status	string	query	False	Filter by statistics.cloud.status • Introduced in: 9.7
statistics.cloud.iops_raw.total	integer	query	False	Filter by statistics.cloud.iops_raw.total • Introduced in: 9.7
statistics.cloud.iops_raw.read	integer	query	False	Filter by statistics.cloud.iops_raw.read • Introduced in: 9.7
statistics.cloud.iops_raw.other	integer	query	False	Filter by statistics.cloud.iops_raw.other • Introduced in: 9.7
statistics.cloud.iops_raw.write	integer	query	False	Filter by statistics.cloud.iops_raw.write • Introduced in: 9.7

Name	Type	In	Required	Description
statistics.cloud.times tamp	string	query	False	Filter by statistics.cloud.times tamp • Introduced in: 9.7
statistics.cloud.laten cy_raw.total	integer	query	False	Filter by statistics.cloud.laten cy_raw.total • Introduced in: 9.7
statistics.cloud.laten cy_raw.read	integer	query	False	Filter by statistics.cloud.laten cy_raw.read • Introduced in: 9.7
statistics.cloud.laten cy_raw.other	integer	query	False	Filter by statistics.cloud.laten cy_raw.other • Introduced in: 9.7
statistics.cloud.laten cy_raw.write	integer	query	False	Filter by statistics.cloud.laten cy_raw.write • Introduced in: 9.7
statistics.throughput _raw.total	integer	query	False	Filter by statistics.throughput _raw.total
statistics.throughput _raw.read	integer	query	False	Filter by statistics.throughput _raw.read
statistics.throughput _raw.other	integer	query	False	Filter by statistics.throughput _raw.other

Name	Type	In	Required	Description
statistics.throughput_raw.write	integer	query	False	Filter by statistics.throughput_raw.write
statistics.flexcache_raw.status	string	query	False	Filter by statistics.flexcache_raw.status • Introduced in: 9.8
statistics.flexcache_raw.cache_miss_blocks	integer	query	False	Filter by statistics.flexcache_raw.cache_miss_blocks • Introduced in: 9.8
statistics.flexcache_raw.client_requested_blocks	integer	query	False	Filter by statistics.flexcache_raw.client_requested_blocks • Introduced in: 9.8
statistics.flexcache_raw.timestamp	string	query	False	Filter by statistics.flexcache_raw.timestamp • Introduced in: 9.8
access_time_enabled	boolean	query	False	Filter by access_time_enabled • Introduced in: 9.8
guarantee.honored	boolean	query	False	Filter by guarantee.honored
guarantee.type	string	query	False	Filter by guarantee.type
encryption.state	string	query	False	Filter by encryption.state

Name	Type	In	Required	Description
encryption.type	string	query	False	Filter by encryption.type
encryption.enabled	boolean	query	False	Filter by encryption.enabled
encryption.key_id	string	query	False	Filter by encryption.key_id
encryption.rekey	boolean	query	False	Filter by encryption.rekey
encryption.status.code	string	query	False	Filter by encryption.status.code
encryption.status.message	string	query	False	Filter by encryption.status.message
analytics.supported	boolean	query	False	Filter by analytics.supported <ul style="list-style-type: none"> Introduced in: 9.8
analytics.state	string	query	False	Filter by analytics.state <ul style="list-style-type: none"> Introduced in: 9.8
analytics.scan_progress	integer	query	False	Filter by analytics.scan_progress <ul style="list-style-type: none"> Introduced in: 9.8
analytics.unsupported_reason.code	string	query	False	Filter by analytics.unsupported_reason.code <ul style="list-style-type: none"> Introduced in: 9.8

Name	Type	In	Required	Description
analytics.unsupported_reason.message	string	query	False	Filter by analytics.unsupported_reason.message • Introduced in: 9.8
comment	string	query	False	Filter by comment
style	string	query	False	Filter by style
space.used	integer	query	False	Filter by space.used
space.total_footprint	integer	query	False	Filter by space.total_footprint • Introduced in: 9.8
space.footprint	integer	query	False	Filter by space.footprint
space.block_storage_inactive_user_data	integer	query	False	Filter by space.block_storage_inactive_user_data
space.performance_tier_footprint	integer	query	False	Filter by space.performance_tier_footprint • Introduced in: 9.8
space.local_tier_footprint	integer	query	False	Filter by space.local_tier_footprint • Introduced in: 9.8
space.metadata	integer	query	False	Filter by space.metadata
space.available	integer	query	False	Filter by space.available

Name	Type	In	Required	Description
space.snapshot.reserve_percent	integer	query	False	Filter by space.snapshot.reserve_percent
space.snapshot.used	integer	query	False	Filter by space.snapshot.used
space.logical_space.enforcement	boolean	query	False	Filter by space.logical_space.enforcement
space.logical_space.used_by_afs	integer	query	False	Filter by space.logical_space.used_by_afs
space.logical_space.available	integer	query	False	Filter by space.logical_space.available
space.logical_space.reporting	boolean	query	False	Filter by space.logical_space.reporting
space.capacity_tier_footprint	integer	query	False	Filter by space.capacity_tier_footprint
space.over_provisioned	integer	query	False	Filter by space.over_provisioned
space.size	integer	query	False	Filter by space.size
tiering.object_tags	string	query	False	Filter by tiering.object_tags • Introduced in: 9.8
tiering.min_cooling_days	integer	query	False	Filter by tiering.min_cooling_days • Introduced in: 9.8

Name	Type	In	Required	Description
tiering.policy	string	query	False	Filter by tiering.policy
type	string	query	False	Filter by type
clone.parent_snapshot.uuid	string	query	False	Filter by clone.parent_snapshot.uuid
clone.parent_snapshot.name	string	query	False	Filter by clone.parent_snapshot.name
clone.split_estimate	integer	query	False	Filter by clone.split_estimate
clone.split_initiated	boolean	query	False	Filter by clone.split_initiated
clone.parent_volume.uuid	string	query	False	Filter by clone.parent_volume.uuid
clone.parent_volume.name	string	query	False	Filter by clone.parent_volume.name
clone.split_complete_percent	integer	query	False	Filter by clone.split_complete_percent
clone.parent_svm.uuid	string	query	False	Filter by clone.parent_svm.uuid
clone.parent_svm.name	string	query	False	Filter by clone.parent_svm.name
clone.is_flexclone	boolean	query	False	Filter by clone.is_flexclone
snapshot_policy.name	string	query	False	Filter by snapshot_policy.name

Name	Type	In	Required	Description
snapshot_policy.uuid	string	query	False	Filter by snapshot_policy.uuid
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
movement.percent_complete	integer	query	False	Filter by movement.percent_complete
movement.destination_aggregate.uuid	string	query	False	Filter by movement.destination_aggregate.uuid
movement.destination_aggregate.name	string	query	False	Filter by movement.destination_aggregate.name
movement.state	string	query	False	Filter by movement.state
movement.cutover_window	integer	query	False	Filter by movement.cutover_window
snaplock.litigation_count	integer	query	False	Filter by snaplock.litigation_count
snaplock.autocommit_period	string	query	False	Filter by snaplock.autocommit_period
snaplock.expiry_time	string	query	False	Filter by snaplock.expiry_time
snaplock.append_mode_enabled	boolean	query	False	Filter by snaplock.append_mode_enabled

Name	Type	In	Required	Description
snaplock.unspecified_retention_file_count	integer	query	False	Filter by snaplock.unspecified_retention_file_count • Introduced in: 9.8
snaplock.is_audit_log	boolean	query	False	Filter by snaplock.is_audit_log
snaplock.privileged_delete	string	query	False	Filter by snaplock.privileged_delete
snaplock.retention.minimum	string	query	False	Filter by snaplock.retention.minimum
snaplock.retention.maximum	string	query	False	Filter by snaplock.retention.maximum
snaplock.retention.default	string	query	False	Filter by snaplock.retention.default
snaplock.compliance_clock_time	string	query	False	Filter by snaplock.compliance_clock_time
snaplock.type	string	query	False	Filter by snaplock.type
cloud_retrieval_policy	string	query	False	Filter by cloud_retrieval_policy • Introduced in: 9.8
autosize.grow_threshold	integer	query	False	Filter by autosize.grow_threshold

Name	Type	In	Required	Description
autosize.shrink_threshold	integer	query	False	Filter by autosize.shrink_threshold
autosize.minimum	integer	query	False	Filter by autosize.minimum
autosize.mode	string	query	False	Filter by autosize.mode
autosize.maximum	integer	query	False	Filter by autosize.maximum
snapmirror.is_protected	boolean	query	False	Filter by snapmirror.is_protected • Introduced in: 9.7
aggregates.uuid	string	query	False	Filter by aggregates.uuid
aggregates.name	string	query	False	Filter by aggregates.name
create_time	string	query	False	Filter by create_time
consistency_group.name	string	query	False	Filter by consistency_group.name • Introduced in: 9.7
nas.security_style	string	query	False	Filter by nas.security_style
nas.path	string	query	False	Filter by nas.path
nas.uid	integer	query	False	Filter by nas.uid
nas.export_policy.id	integer	query	False	Filter by nas.export_policy.id

Name	Type	In	Required	Description
nas.export_policy.name	string	query	False	Filter by nas.export_policy.name
nas.gid	integer	query	False	Filter by nas.gid
nas.unix_permissions	integer	query	False	Filter by nas.unix_permissions
language	string	query	False	Filter by language
is_svm_root	boolean	query	False	Filter by is_svm_root <ul style="list-style-type: none"> Introduced in: 9.7
flexcache_endpoint_type	string	query	False	Filter by flexcache_endpoint_type
name	string	query	False	Filter by name
metric.latency.total	integer	query	False	Filter by metric.latency.total
metric.latency.read	integer	query	False	Filter by metric.latency.read
metric.latency.other	integer	query	False	Filter by metric.latency.other
metric.latency.write	integer	query	False	Filter by metric.latency.write
metric.cloud.duration	string	query	False	Filter by metric.cloud.duration <ul style="list-style-type: none"> Introduced in: 9.7

Name	Type	In	Required	Description
metric.cloud.timestamp	string	query	False	Filter by metric.cloud.timestamp • Introduced in: 9.7
metric.cloud.latency.total	integer	query	False	Filter by metric.cloud.latency.total • Introduced in: 9.7
metric.cloud.latency.read	integer	query	False	Filter by metric.cloud.latency.read • Introduced in: 9.7
metric.cloud.latency.other	integer	query	False	Filter by metric.cloud.latency.other • Introduced in: 9.7
metric.cloud.latency.write	integer	query	False	Filter by metric.cloud.latency.write • Introduced in: 9.7
metric.cloud.iops.total	integer	query	False	Filter by metric.cloud.iops.total • Introduced in: 9.7
metric.cloud.iops.read	integer	query	False	Filter by metric.cloud.iops.read • Introduced in: 9.7

Name	Type	In	Required	Description
metric.cloud.iops.other	integer	query	False	Filter by metric.cloud.iops.other • Introduced in: 9.7
metric.cloud.iops.write	integer	query	False	Filter by metric.cloud.iops.write • Introduced in: 9.7
metric.cloud.status	string	query	False	Filter by metric.cloud.status • Introduced in: 9.7
metric.flexcache.timestamp	string	query	False	Filter by metric.flexcache.timestamp • Introduced in: 9.8
metric.flexcache.cache_miss_percent	integer	query	False	Filter by metric.flexcache.cache_miss_percent • Introduced in: 9.8
metric.flexcache.duration	string	query	False	Filter by metric.flexcache.duration • Introduced in: 9.8
metric.flexcache.status	string	query	False	Filter by metric.flexcache.status • Introduced in: 9.8

Name	Type	In	Required	Description
metric.status	string	query	False	Filter by metric.status
metric.timestamp	string	query	False	Filter by metric.timestamp
metric.duration	string	query	False	Filter by metric.duration
metric.throughput.total	integer	query	False	Filter by metric.throughput.total
metric.throughput.read	integer	query	False	Filter by metric.throughput.read
metric.throughput.other	integer	query	False	Filter by metric.throughput.other
metric.throughput.write	integer	query	False	Filter by metric.throughput.write
metric.iops.total	integer	query	False	Filter by metric.iops.total
metric.iops.read	integer	query	False	Filter by metric.iops.read
metric.iops.other	integer	query	False	Filter by metric.iops.other
metric.iops.write	integer	query	False	Filter by metric.iops.write
uuid	string	query	False	Filter by uuid
qos.policy.max_throughput_mbps	integer	query	False	Filter by qos.policy.max_throughput_mbps
qos.policy.max_throughput_iops	integer	query	False	Filter by qos.policy.max_throughput_iops

Name	Type	In	Required	Description
qos.policy.uuid	string	query	False	Filter by qos.policy.uuid
qos.policy.min_throu ghput_iops	integer	query	False	Filter by qos.policy.min_throu ghput_iops
qos.policy.min_throu ghput_mbps	integer	query	False	Filter by qos.policy.min_throu ghput_mbps • Introduced in: 9.8
qos.policy.name	string	query	False	Filter by qos.policy.name
size	integer	query	False	Filter by size
error_state.has_bad _blocks	boolean	query	False	Filter by error_state.has_bad _blocks
error_state.is_incons istent	boolean	query	False	Filter by error_state.is_incon sistent
quota.state	string	query	False	Filter by quota.state
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[volume]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "aggregates": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "aggr1",
          "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
        }
      ],
      "analytics": {
        "scan_progress": "17",
        "state": "string",
        "unsupported_reason": {
          "code": "111411207",
          "message": "File system analytics cannot be enabled on
volumes that contain LUNs."
        }
      },
      "application": {
        "name": "string",
        "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
      },
      "autosize": {
        "mode": "string"
      },
      "clone": {
```

```

"parent_snapshot": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "this_snapshot",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"parent_svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"parent_volume": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "volume1",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
},
"split_complete_percent": 0,
"split_estimate": 0
},
"cloud_retrieval_policy": "string",
"comment": "string",
"consistency_group": {
  "name": "consistency_group_1"
},
"create_time": "2018-06-04T19:00:00Z",
"efficiency": {
  "compaction": "string",
  "compression": "string",
  "cross_volume_dedupe": "string",
  "dedupe": "string",
  "policy": {
    "name": "string"
  }
},
"encryption": {

```

```

    "key_id": "string",
    "state": "string",
    "status": {
        "code": "string",
        "message": "string"
    },
    "type": "string"
},
"files": {
    "used": 0
},
"flexcache_endpoint_type": "string",
"guarantee": {
    "type": "string"
},
"language": "string",
"metric": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    }
},
"cloud": {
    "duration": "PT15S",
    "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
},
"duration": "PT15S",
"flexcache": {
    "cache_miss_percent": "20",
    "duration": "PT1D",
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
},
"iops": {
    "read": "200",

```

```

        "total": "1000",
        "write": "100"
    },
    "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "status": "ok",
    "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"movement": {
    "cutover_window": "30",
    "destination_aggregate": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "state": "replicating",
    "tiering_policy": "string"
},
"name": "vol_cs_dept",
"nas": {
    "export_policy": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "id": "100",
        "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": "0755"
},

```

```

"qos": {
  "policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
},
"quota": {
  "state": "string"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04T19:00:00Z",
  "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": "10",
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise",
  "unspecified_retention_file_count": "10"
},
"snapshot_policy": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,

```



```

"capacity_tier_footprint": 0,
"footprint": 0,
"local_tier_footprint": 0,
"logical_space": {
  "available": 0,
  "used_by_afs": 0
},
"metadata": 0,
"over_provisioned": 0,
"performance_tier_footprint": 0,
"snapshot": {
  "used": 0
},
"total_footprint": 0,
"used": 0
},
"state": "string",
"statistics": {
  "cloud": {
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "flexcache_raw": {
    "cache_miss_blocks": "10",
    "client_requested_blocks": "500",
    "status": "ok",
    "timestamp": "2017-01-25T11:20:13Z"
  },
  "iops_raw": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency_raw": {
    "read": "200",
    "total": "1000",

```

```

        "write": "100"
      },
      "status": "ok",
      "throughput_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    },
    "style": "string",
    "svm": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "tiering": {
      "object_tags": [
        "string"
      ],
      "policy": "string",
      "supported": null
    },
    "type": "string",
    "use_mirrored_aggregates": null,
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

unsupported_reason

Name	Type	Description
code	string	If file system analytics is not supported on the volume, this field provides the error code explaining why.
message	string	If file system analytics is not supported on the volume, this field provides the error message explaining why.

analytics

Name	Type	Description
scan_progress	integer	Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is initializing.
state	string	<p>File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information all existing files and directories is currently running. If this value is 'unknown' that means there was an internal error when determining the file system analytics state for the volume.</p> <ul style="list-style-type: none"> enum: ["unknown", "initializing", "off", "on"] Introduced in: 9.8
supported	boolean	This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field.
unsupported_reason	unsupported_reason	

application

Name	Type	Description
name	string	Name of the application to which the volume belongs. Available only when the volume is part of an application.

Name	Type	Description
uuid	string	UUID of the application to which the volume belongs. Available only when the volume is part of an application.

autosize

Name	Type	Description
grow_threshold	integer	Used space threshold size, in percentage, for the automatic growth of the volume. When the amount of used space in the volume becomes greater than this threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size..
maximum	integer	Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume.
minimum	integer	Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume.
mode	string	Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled.

Name	Type	Description
shrink_threshold	integer	Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size.

snapshot_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parent_svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

parent_volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

clone

Name	Type	Description
is_flexclone	boolean	Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.
split_initiated	boolean	This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.

consistency_group

Consistency group the volume is part of.

Name	Type	Description
name	string	Name of the consistency group.

policy

Name	Type	Description
name	string	Specifies the name of the efficiency policy.

efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.

Name	Type	Description
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.
dedupe	string	The system can be enabled/disabled dedupe. inline ‐ Data will be deduped first and written to the volume. background ‐ Data will be written to the volume and deduped later. both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.
policy	policy	

status

Name	Type	Description
code	string	Encryption progress message code.
message	string	Encryption progress message.

encryption

Name	Type	Description
enabled	boolean	Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate.
key_id	string	The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key.
rekey	boolean	If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	

Name	Type	Description
type	string	Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption). aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption).

error_state

Name	Type	Description
has_bad_blocks	boolean	Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned.
is_inconsistent	boolean	Indicates whether the file system has any inconsistencies. true ‐ File system is inconsistent. false ‐ File system is not inconsistent.

files

Name	Type	Description
maximum	integer	The maximum number of files (inodes) for user-visible data allowed on the volume. This value can be increased or decreased. Increasing the maximum number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and this value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume. Valid in PATCH.
used	integer	Number of files (inodes) used for user-visible data permitted on the volume. This field is valid only when the volume is online.

guarantee

Name	Type	Description
honored	boolean	Is the space guarantee of this volume honored in the aggregate?
type	string	The type of space guarantee of this volume in the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.

Name	Type	Description
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

flexcache

Performance number for FlexCache used to measure cache effectiveness.

Name	Type	Description
cache_miss_percent	integer	Cache miss percentage.
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS, latency and throughput.

Name	Type	Description
<code>_links</code>	_links	
<code>cloud</code>	cloud	Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.
<code>duration</code>	<code>string</code>	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
<code>flexcache</code>	flexcache	Performance number for FlexCache used to measure cache effectiveness.
<code>iops</code>	iops	The rate of I/O operations observed at the storage object.
<code>latency</code>	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

destination_aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
destination_aggregate	destination_aggregate	Aggregate
percent_complete	integer	Completion percentage
state	string	State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress.
tiering_policy	string	Tiering policy for FabricPool

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	export_policy	Export Policy

Name	Type	Description
gid	integer	The UNIX group ID of the volume. Valid in POST or PATCH.
path	string	The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline removes its junction path. This attribute is reported in GET only when the volume is mounted.
security_style	string	Security style associated with the volume. Valid in POST or PATCH. mixed ‐ Mixed-style security ntfs ‐ NTFS/Windows-style security unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions unix ‐ Unix-style security.
uid	integer	The UNIX user ID of the volume. Valid in POST or PATCH.

Name	Type	Description
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero.

policy

When "min_throughput_iops", "min_throughput_mbps", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

Name	Type	Description
_links	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

qos

QoS information

Name	Type	Description
policy	policy	When "min_throughput_iops", "min_throughput_mbps", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

Name	Type	Description
enabled	boolean	This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off".
state	string	Quota state of the volume

retention

Name	Type	Description
default	string	<p>Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num></p>

Name	Type	Description
maximum	string	Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.

Name	Type	Description
minimum	string	Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num>

snaplock

Name	Type	Description
append_mode_enabled	boolean	Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires.

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.

Name	Type	Description
privileged_delete	string	Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled.
retention	retention	
type	string	The SnapLock type of the volume. compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files. enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume. non_snaplock ‐ Indicates the volume is non-snaplock.
unspecified_retention_file_count	integer	Indicates the number of files with an unspecified retention time in the volume.

snapmirror

Specifies attributes for SnapMirror protection.

Name	Type	Description
is_protected	boolean	Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data.

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.
enforcement	boolean	Specifies whether space accounting for operations on the volume is done along with storage efficiency.
reporting	boolean	Specifies whether space reporting on the volume is done along with storage efficiency.
used_by_afs	integer	The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.

snapshot

Name	Type	Description
autodelete_enabled	boolean	Specifies whether Snapshot copy autodelete is currently enabled on this volume.
reserve_percent	integer	The space that has been set aside as a reserve for Snapshot copy usage, in percent.
used	integer	The total space used by Snapshot copies in the volume, in bytes.

space

Name	Type	Description
available	integer	The available space, in bytes.
block_storage_inactive_user_data	integer	The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store.
capacity_tier_footprint	integer	Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.
footprint	integer	Data used for this volume in the aggregate, in bytes.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
metadata	integer	Space used by the volume metadata in the aggregate, in bytes.
over_provisioned	integer	The amount of space not available for this volume in the aggregate, in bytes.
performance_tier_footprint	integer	Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.
size	integer	Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.
snapshot	snapshot	
total_footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
used	integer	The virtual space used (includes volume reserves) before storage efficiency, in bytes.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

Name	Type	Description
cache_miss_blocks	integer	Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage.
client_requested_blocks	integer	Total blocks requested by the client.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
cloud	cloud	These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.
flexcache_raw	flexcache_raw	Performance numbers for FlexCache used to measure cache effectiveness.
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

Name	Type	Description
min_cooling_days	integer	<p>This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data.</p> <p>This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.</p>
object_tags	array[string]	<p>This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value".</p>

Name	Type	Description
policy	string	Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks. auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store none ‐ Volume blocks will not be tiered to the cloud store. snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.

volume

Name	Type	Description
_links	_links	
access_time_enabled	boolean	Indicates whether or not access time updates are enabled on the volume.
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
analytics	analytics	
application	application	

Name	Type	Description
autosize	autosize	
clone	clone	
cloud_retrieval_policy	string	<p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved. "on_read" policy retrieves tiered data for all client driven data reads.</p> <p>"never" policy never retrieves tiered data.</p> <p>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p>
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.
create_time	string	Creation time of the volume. This field is generated when the volume is created.
efficiency	efficiency	

Name	Type	Description
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	metric	Performance numbers, such as IOPS, latency and throughput.
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapmirror	snapmirror	Specifies attributes for SnapMirror protection.
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	
state	string	Volume state. A volume can only be brought online if it is offline. Taking a volume offline removes its junction path. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data.

Name	Type	Description
statistics	statistics	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.
style	string	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. flexvol ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a volume on an SVM and storage aggregates

POST /storage/volumes

Introduced In: 9.6

Creates a volume on a specified SVM and storage aggregates.

Required properties

- `svm.uuid` or `svm.name` - Existing SVM in which to create the volume.
- `name` - Name of the volume.
- `aggregates.name` or `aggregates.uuid` - Existing aggregates in which to create the volume.

Default property values

- state - *online*
- size - *20MB*
- style - *flexvol*
- type - *rw*
- encryption.enabled - *false*
- snapshot_policy.name - *default*
- guarantee.type - *volume*

Related ONTAP commands

- volume create
- volume clone create

Parameters

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
analytics	analytics	
application	application	

Name	Type	Description
autosize	autosize	
clone	clone	
cloud_retrieval_policy	string	<p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.</p> <p>"on_read" policy retrieves tiered data for all client driven data reads.</p> <p>"never" policy never retrieves tiered data.</p> <p>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p>
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents_per_aggregate	integer	Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup. If a volume is being created on a single aggregate, the system will create a flexible volume if the "constituents_per_aggregate" field is not specified, and a FlexGroup if it is specified. If a volume is being created on multiple aggregates, the system will always create a FlexGroup.
create_time	string	Creation time of the volume. This field is generated when the volume is created.
efficiency	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

Name	Type	Description
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapmirror	snapmirror	Specifies attributes for SnapMirror protection.
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. Taking a volume offline removes its junction path. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data.
style	string	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. flexvol ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6

Example request

```
{
  "access_time_enabled": null,
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "analytics": {
    "scan_progress": "17",
    "state": "string",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "autosize": {
    "mode": "string"
  },
  "clone": {
    "parent_snapshot": {
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "cloud_retrieval_policy": "string",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1"
  }
}
```

```

},
"create_time": "2018-06-04T19:00:00Z",
"efficiency": {
  "compaction": "string",
  "compression": "string",
  "cross_volume_dedupe": "string",
  "dedupe": "string",
  "policy": {
    "name": "string"
  }
},
},
"encryption": {
  "key_id": "string",
  "rekey": null,
  "state": "string",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "string"
},
"files": {
  "used": 0
},
"flexcache_endpoint_type": "string",
"guarantee": {
  "type": "string"
},
"language": "string",
"movement": {
  "cutover_window": "30",
  "destination_aggregate": {
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": 0,
  "state": "replicating",
  "tiering_policy": "string"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "id": "100",
    "name": "default"
  },
  "path": "/user/my_volume",

```

```

    "security_style": "string",
    "unix_permissions": "0755"
  },
  "qos": {
    "policy": {
      "max_throughput_iops": "10000",
      "max_throughput_mbps": "500",
      "min_throughput_iops": "2000",
      "min_throughput_mbps": "500",
      "name": "performance",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  },
  "quota": {
    "state": "string"
  },
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04T19:00:00Z",
    "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
    "is_audit_log": 1,
    "litigation_count": "10",
    "privileged_delete": "enabled",
    "retention": {
      "default": "P30Y",
      "maximum": "P30Y",
      "minimum": "P30Y"
    },
    "type": "enterprise",
    "unspecified_retention_file_count": "10"
  },
  "snapshot_policy": {
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "space": {
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "capacity_tier_footprint": 0,
    "footprint": 0,
    "local_tier_footprint": 0,
    "logical_space": {
      "available": 0,
      "used_by_afs": 0
    }
  },

```

```

    "metadata": 0,
    "over_provisioned": 0,
    "performance_tier_footprint": 0,
    "snapshot": {
      "used": 0
    },
    "total_footprint": 0,
    "used": 0
  },
  "state": "string",
  "style": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "object_tags": [
      "string"
    ],
    "policy": "string"
  },
  "type": "string",
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

{
  "job": {
    "uuid": "string"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
787140	One of "aggregates.uuid", "aggregates.name", or "style" must be provided.
787141	The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates.
917526	The volume name specified is a duplicate.
917829	Volume autosize grow threshold must be larger than autosize shrink threshold.
917831	Volume minimum autosize must be smaller than the maximum autosize.
917835	Maximum allowed snapshot.reserve_percent value during a volume creation is 90. Use PATCH to set it to a higher value after the volume has been created.
918191	Flexvol tiering min cooling days requires an effective cluster version of ONTAP 9.4 or later.
918194	Tiering min cooling days not supported for SVMDR.
918195	Tiering min cooling days not supported for non data volumes.
918196	Tiering min cooling days not allowed for the provided tiering policy.
918215	FlexGroup tiering min cooling days requires an effective cluster version of ONTAP 9.5 or later.
918233	The target field cannot be specified for this operation.
918236	The specified "parent_volume.uuid" and "parent_volume.name" do not refer to the same volume.
918240	The target style is an invalid volume style.
918241	The target style is an unsupported volume style for volume creation.
918242	When creating a flexible volume, exactly one aggregate must be specified via either "aggregates.name" or "aggregates.uuid".
918243	The specified Snapshot copy UUID is not correct for the specified Snapshot copy name.
918244	Invalid "volume.type" for clone volume.

Error Code	Description
918246	"volume.clone.parent_volume.name" or "volume.clone.parent_volume.uuid" must be provided.
918247	Specifying a value is not valid for a volume FlexClone creation.
918252	"nas.path" is invalid.
918290	cloud retrieval policy requires an effective cluster version of 9.8 or later.
918291	Invalid volume cloud retrieval policy for the provided tiering policy.
918292	cloud retrieval policy not supported for non data volume.
918521	The volume maximum autosize must be smaller than or equal to the maximum volume size.
918524	Volume minimum autosize must be less than or equal to the current volume size.
2621706	The specified "svm.uuid" and "svm.name" do not refer to the same SVM.
2621707	No SVM was specified. Either "svm.name" or "svm.uuid" must be supplied.
111411205	File system analytics requires an effective cluster version of 9.8 or later.
111411206	The specified "analytics.state" is invalid.
111411207	File system analytics cannot be enabled on volumes that contain LUNs.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

aggregates

Aggregate

Name	Type	Description
name	string	
uuid	string	

unsupported_reason

Name	Type	Description
code	string	If file system analytics is not supported on the volume, this field provides the error code explaining why.
message	string	If file system analytics is not supported on the volume, this field provides the error message explaining why.

analytics

Name	Type	Description
scan_progress	integer	Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is <code>initializing</code> .

Name	Type	Description
state	string	<p>File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information all existing files and directories is currently running. If this value is 'unknown' that means there was an internal error when determining the file system analytics state for the volume.</p> <ul style="list-style-type: none"> enum: ["unknown", "initializing", "off", "on"] Introduced in: 9.8
supported	boolean	This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field.
unsupported_reason	unsupported_reason	

application

Name	Type	Description
name	string	Name of the application to which the volume belongs. Available only when the volume is part of an application.
uuid	string	UUID of the application to which the volume belongs. Available only when the volume is part of an application.

autosize

Name	Type	Description
grow_threshold	integer	Used space threshold size, in percentage, for the automatic growth of the volume. When the amount of used space in the volume becomes greater than this threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size..
maximum	integer	Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume.
minimum	integer	Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume.
mode	string	Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled.

Name	Type	Description
shrink_threshold	integer	Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size.

snapshot_reference

Name	Type	Description
name	string	
uuid	string	

parent_svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

parent_volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

clone

Name	Type	Description
is_flexclone	boolean	Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.
split_initiated	boolean	This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone from FlexVol. Valid in PATCH.

consistency_group

Consistency group the volume is part of.

Name	Type	Description
name	string	Name of the consistency group.

policy

Name	Type	Description
name	string	Specifies the name of the efficiency policy.

efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.

Name	Type	Description
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.
dedupe	string	The system can be enabled/disabled dedupe. inline ‐ Data will be deduped first and written to the volume. background ‐ Data will be written to the volume and deduped later. both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.
policy	policy	

status

Name	Type	Description
code	string	Encryption progress message code.
message	string	Encryption progress message.

encryption

Name	Type	Description
enabled	boolean	Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate.
key_id	string	The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key.
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	
type	string	Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption). aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption).

error_state

Name	Type	Description
has_bad_blocks	boolean	Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned.
is_inconsistent	boolean	Indicates whether the file system has any inconsistencies. true ‐ File system is inconsistent. false ‐ File system is not inconsistent.

files

Name	Type	Description
maximum	integer	The maximum number of files (inodes) for user-visible data allowed on the volume. This value can be increased or decreased. Increasing the maximum number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and this value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume. Valid in PATCH.
used	integer	Number of files (inodes) used for user-visible data permitted on the volume. This field is valid only when the volume is online.

guarantee

Name	Type	Description
honored	boolean	Is the space guarantee of this volume honored in the aggregate?
type	string	The type of space guarantee of this volume in the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.

Name	Type	Description
write	integer	Peformance metric for write I/O operations.

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.

Name	Type	Description
timestamp	string	The timestamp of the performance data.

flexcache

Performance number for FlexCache used to measure cache effectiveness.

Name	Type	Description
cache_miss_percent	integer	Cache miss percentage.
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS, latency and throughput.

Name	Type	Description
cloud	cloud	Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
flexcache	flexcache	Performance number for FlexCache used to measure cache effectiveness.
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

destination_aggregate

Aggregate

Name	Type	Description
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
destination_aggregate	destination_aggregate	Aggregate
percent_complete	integer	Completion percentage
state	string	State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress.
tiering_policy	string	Tiering policy for FabricPool

export_policy

Export Policy

Name	Type	Description
id	integer	
name	string	

nas

Name	Type	Description
export_policy	export_policy	Export Policy
gid	integer	The UNIX group ID of the volume. Valid in POST or PATCH.

Name	Type	Description
path	string	The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline removes its junction path. This attribute is reported in GET only when the volume is mounted.
security_style	string	Security style associated with the volume. Valid in POST or PATCH. mixed ‐ Mixed-style security ntfs ‐ NTFS/Windows-style security unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions unix ‐ Unix-style security.
uid	integer	The UNIX user ID of the volume. Valid in POST or PATCH.

Name	Type	Description
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero.

policy

When "min_throughput_iops", "min_throughput_mbps", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

Name	Type	Description
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

qos

QoS information

Name	Type	Description
policy	policy	When "min_throughput_iops", "min_throughput_mbps", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

Name	Type	Description
enabled	boolean	This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off".
state	string	Quota state of the volume

retention

Name	Type	Description
default	string	<p>Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num></p>

Name	Type	Description
maximum	string	<p>Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</p>

Name	Type	Description
minimum	string	Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num>

snaplock

Name	Type	Description
append_mode_enabled	boolean	Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires.

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.

Name	Type	Description
privileged_delete	string	Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled.
retention	retention	
type	string	The SnapLock type of the volume. compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files. enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume. non_snaplock ‐ Indicates the volume is non-snaplock.
unspecified_retention_file_count	integer	Indicates the number of files with an unspecified retention time in the volume.

snapmirror

Specifies attributes for SnapMirror protection.

Name	Type	Description
is_protected	boolean	Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data.

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.
enforcement	boolean	Specifies whether space accounting for operations on the volume is done along with storage efficiency.
reporting	boolean	Specifies whether space reporting on the volume is done along with storage efficiency.
used_by_afs	integer	The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.

snapshot

Name	Type	Description
autodelete_enabled	boolean	Specifies whether Snapshot copy autodelete is currently enabled on this volume.
reserve_percent	integer	The space that has been set aside as a reserve for Snapshot copy usage, in percent.
used	integer	The total space used by Snapshot copies in the volume, in bytes.

space

Name	Type	Description
available	integer	The available space, in bytes.

Name	Type	Description
block_storage_inactive_user_data	integer	The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store.
capacity_tier_footprint	integer	Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.
footprint	integer	Data used for this volume in the aggregate, in bytes.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
metadata	integer	Space used by the volume metadata in the aggregate, in bytes.
over_provisioned	integer	The amount of space not available for this volume in the aggregate, in bytes.
performance_tier_footprint	integer	Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.
size	integer	Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.
snapshot	snapshot	
total_footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
used	integer	The virtual space used (includes volume reserves) before storage efficiency, in bytes.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

Name	Type	Description
cache_miss_blocks	integer	Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage.
client_requested_blocks	integer	Total blocks requested by the client.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
cloud	cloud	These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.
flexcache_raw	flexcache_raw	Performance numbers for FlexCache used to measure cache effectiveness.
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

Name	Type	Description
min_cooling_days	integer	<p>This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data.</p> <p>This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.</p>
object_tags	array[string]	<p>This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value".</p>

Name	Type	Description
policy	string	<p>Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH.</p> <p>all &dash; This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks.</p> <p>auto &dash; This policy allows tiering of both snapshot and active file system user data to the cloud store</p> <p>none &dash; Volume blocks will not be tiered to the cloud store.</p> <p>snapshot_only &dash; This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.</p>

Name	Type	Description
supported	boolean	This parameter specifies whether or not FabricPools are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only FabricPool aggregates are used if this parameter is set to true and only non FabricPool aggregates are used if this parameter is set to false. Tiering support for a FlexGroup can be changed by moving all of the constituents to the required aggregates. Note that in order to tier data, not only does the volume need to support tiering by using FabricPools, the tiering "policy" must not be 'none'. A volume that uses FabricPools but has a tiering "policy" of 'none' supports tiering, but will not tier any data.

volume

Name	Type	Description
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
analytics	analytics	
application	application	
autosize	autosize	
clone	clone	

Name	Type	Description
cloud_retrieval_policy	string	<p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved. "on_read" policy retrieves tiered data for all client driven data reads.</p> <p>"never" policy never retrieves tiered data.</p> <p>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p>
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents_per_aggregate	integer	Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup. If a volume is being created on a single aggregate, the system will create a flexible volume if the "constituents_per_aggregate" field is not specified, and a FlexGroup if it is specified. If a volume is being created on multiple aggregates, the system will always create a FlexGroup.
create_time	string	Creation time of the volume. This field is generated when the volume is created.
efficiency	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.

Name	Type	Description
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapmirror	snapmirror	Specifies attributes for SnapMirror protection.
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. Taking a volume offline removes its junction path. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data.
style	string	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. flexvol ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.

Name	Type	Description
use_mirrored_aggregates	boolean	Specifies whether mirrored aggregates are selected when provisioning a FlexGroup without specifying "aggregates.name" or "aggregates.uuid". Only mirrored aggregates are used if this parameter is set to 'true' and only unmirrored aggregates are used if this parameter is set to 'false'. Aggregate level mirroring for a FlexGroup can be changed by moving all of the constituents to the required aggregates. The default value is 'true' for a MetroCluster configuration and is 'false' for a non-MetroCluster configuration.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a volume

DELETE /storage/volumes/{uuid}

Introduced In: 9.6

Deletes a volume. If the UUID belongs to a volume, all of its blocks are freed and returned to its containing aggregate. If a volume is online, it is offlined before deletion. If a volume is mounted, unmount the volume by specifying the nas.path as empty before deleting it using the DELETE operation.

Related ONTAP commands

- `volume delete`
- `volume clone delete`

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of the volume.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve a volume

GET /storage/volumes/{uuid}

Introduced In: 9.6

Retrieves a volume. The GET API can be used to retrieve the quota state for a FlexVol or a FlexGroup volume.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `is_svm_root`
- `analytics.*`
- `application.*`
- `encryption.*`
- `clone.parent_snapshot.name`
- `clone.parent_snapshot.uuid`
- `clone.parent_svm.name`
- `clone.parent_svm.uuid`
- `clone.parent_volume.name`
- `clone.parent_volume.uuid`
- `clone.split_complete_percent`
- `clone.split_estimate`
- `clone.split_initiated`
- `efficiency.*`
- `error_state.*`
- `files.*`
- `nas.export_policy.id`
- `nas.gid`
- `nas.path`
- `nas.security_style`
- `nas.uid`
- `nas.unix_permissions`
- `snaplock.*`
- `restore_to.*`
- `snapshot_policy.uuid`
- `quota.*`
- `qos.*`
- `flexcache_endpoint_type`

- space.block_storage_inactive_user_data
- space.capacity_tier_footprint
- space.performance_tier_footprint
- space.local_tier_footprint
- space.footprint
- space.over_provisioned
- space.metadata
- space.total_footprint
- space.logical_space.*
- space.snapshot.*
- guarantee.*
- autosize.*
- movement.*
- statistics.*

Related ONTAP commands

- volume show
- volume clone show
- volume efficiency show
- volume encryption show
- volume flexcache show
- volume flexgroup show
- volume move show
- volume quota show
- volume show-space
- volume snaplock show

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of the volume.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
access_time_enabled	boolean	Indicates whether or not access time updates are enabled on the volume.
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
analytics	analytics	
application	application	
autosize	autosize	
clone	clone	

Name	Type	Description
cloud_retrieval_policy	string	<p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.</p> <p>"on_read" policy retrieves tiered data for all client driven data reads.</p> <p>"never" policy never retrieves tiered data.</p> <p>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p>
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.
create_time	string	Creation time of the volume. This field is generated when the volume is created.
efficiency	efficiency	
encryption	encryption	
error_state	error_state	
files	files	

Name	Type	Description
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.
language	string	Language encoding setting for volume. If no language is specified, the volume inherits its SVM language encoding setting.
metric	metric	Performance numbers, such as IOPS, latency and throughput.
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information

Name	Type	Description
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapmirror	snapmirror	Specifies attributes for SnapMirror protection.
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	
state	string	Volume state. A volume can only be brought online if it is offline. Taking a volume offline removes its junction path. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data.
statistics	statistics	These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
style	string	The style of the volume. If "style" is not specified, the volume type is determined based on the specified aggregates. Specifying a single aggregate, without "constituents_per_aggregate", creates a flexible volume. Specifying multiple aggregates, or a single aggregate with "constituents_per_aggregate", creates a FlexGroup. Specifying a volume "style" creates a volume of that type. For example, if the style is "flexvol" you must specify a single aggregate. If the style is "flexgroup", the system either uses the specified aggregates or automatically provisions aggregates if there are no specified aggregates. flexvol ‐ flexible volumes and FlexClone volumes flexgroup ‐ FlexGroups.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
type	string	Type of the volume. rw ‐ read-write volume. dp ‐ data-protection volume. ls ‐ load-sharing <code>dp</code> volume. Valid in GET.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "aggregates": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "analytics": {
    "scan_progress": "17",
    "state": "string",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "autosize": {
    "mode": "string"
  },
  "clone": {
    "parent_snapshot": {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
```

```

    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "parent_volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  },
  "split_complete_percent": 0,
  "split_estimate": 0
},
"cloud_retrieval_policy": "string",
"comment": "string",
"consistency_group": {
  "name": "consistency_group_1"
},
"create_time": "2018-06-04T19:00:00Z",
"efficiency": {
  "compaction": "string",
  "compression": "string",
  "cross_volume_dedupe": "string",
  "dedupe": "string",
  "policy": {
    "name": "string"
  }
},
"encryption": {
  "key_id": "string",
  "state": "string",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "string"
},
"files": {
  "used": 0
}

```

```

},
"flexcache_endpoint_type": "string",
"guarantee": {
  "type": "string"
},
"language": "string",
"metric": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
},
"cloud": {
  "duration": "PT15S",
  "iops": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "latency": {
    "read": "200",
    "total": "1000",
    "write": "100"
  },
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"duration": "PT15S",
"flexcache": {
  "cache_miss_percent": "20",
  "duration": "PT1D",
  "status": "ok",
  "timestamp": "2017-01-25T11:20:13Z"
},
"iops": {
  "read": "200",
  "total": "1000",
  "write": "100"
},
"latency": {
  "read": "200",
  "total": "1000",
  "write": "100"
},
"status": "ok",
"throughput": {

```

```

        "read": "200",
        "total": "1000",
        "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
},
"movement": {
    "cutover_window": "30",
    "destination_aggregate": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "name": "aggr1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "percent_complete": 0,
    "state": "replicating",
    "tiering_policy": "string"
},
"name": "vol_cs_dept",
"nas": {
    "export_policy": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "id": "100",
        "name": "default"
    },
    "path": "/user/my_volume",
    "security_style": "string",
    "unix_permissions": "0755"
},
"qos": {
    "policy": {
        "_links": {
            "self": {
                "href": "/api/resourcelink"
            }
        },
        "max_throughput_iops": "10000",
        "max_throughput_mbps": "500",
        "min_throughput_iops": "2000",

```

```

    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  },
  "quota": {
    "state": "string"
  },
  },
  "snaplock": {
    "append_mode_enabled": "",
    "autocommit_period": "P30M",
    "compliance_clock_time": "2018-06-04T19:00:00Z",
    "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
    "is_audit_log": 1,
    "litigation_count": "10",
    "privileged_delete": "enabled",
    "retention": {
      "default": "P30Y",
      "maximum": "P30Y",
      "minimum": "P30Y"
    },
    "type": "enterprise",
    "unspecified_retention_file_count": "10"
  },
  },
  "snapshot_policy": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "default",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  },
  "space": {
    "available": 0,
    "block_storage_inactive_user_data": 0,
    "capacity_tier_footprint": 0,
    "footprint": 0,
    "local_tier_footprint": 0,
    "logical_space": {
      "available": 0,
      "used_by_afs": 0
    },
    "metadata": 0,
    "over_provisioned": 0,
    "performance_tier_footprint": 0,

```

```

    "snapshot": {
      "used": 0
    },
    "total_footprint": 0,
    "used": 0
  },
  "state": "string",
  "statistics": {
    "cloud": {
      "iops_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency_raw": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "timestamp": "2017-01-25T11:20:13Z"
    },
    "flexcache_raw": {
      "cache_miss_blocks": "10",
      "client_requested_blocks": "500",
      "status": "ok",
      "timestamp": "2017-01-25T11:20:13Z"
    },
    "iops_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "latency_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "status": "ok",
    "throughput_raw": {
      "read": "200",
      "total": "1000",
      "write": "100"
    },
    "timestamp": "2017-01-25T11:20:13Z"
  },
},

```

```

"style": "string",
"svm": {
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "svm1",
  "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
},
"tiering": {
  "object_tags": [
    "string"
  ],
  "policy": "string",
  "supported": null
},
"type": "string",
"use_mirrored_aggregates": null,
"uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

aggregates

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

unsupported_reason

Name	Type	Description
code	string	If file system analytics is not supported on the volume, this field provides the error code explaining why.
message	string	If file system analytics is not supported on the volume, this field provides the error message explaining why.

analytics

Name	Type	Description
scan_progress	integer	Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is <code>initializing</code> .

Name	Type	Description
state	string	<p>File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information all existing files and directories is currently running. If this value is 'unknown' that means there was an internal error when determining the file system analytics state for the volume.</p> <ul style="list-style-type: none"> enum: ["unknown", "initializing", "off", "on"] Introduced in: 9.8
supported	boolean	This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field.
unsupported_reason	unsupported_reason	

application

Name	Type	Description
name	string	Name of the application to which the volume belongs. Available only when the volume is part of an application.
uuid	string	UUID of the application to which the volume belongs. Available only when the volume is part of an application.

autosize

Name	Type	Description
grow_threshold	integer	Used space threshold size, in percentage, for the automatic growth of the volume. When the amount of used space in the volume becomes greater than this threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size..
maximum	integer	Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume.
minimum	integer	Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume.
mode	string	Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled.

Name	Type	Description
shrink_threshold	integer	Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size.

snapshot_reference

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parent_svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

parent_volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.

Name	Type	Description
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6

clone

Name	Type	Description
is_flexclone	boolean	Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.
split_initiated	boolean	This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone form FlexVol. Valid in PATCH.

consistency_group

Consistency group the volume is part of.

Name	Type	Description
name	string	Name of the consistency group.

policy

Name	Type	Description
name	string	Specifies the name of the efficiency policy.

efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.

Name	Type	Description
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.
dedupe	string	The system can be enabled/disabled dedupe. inline ‐ Data will be deduped first and written to the volume. background ‐ Data will be written to the volume and deduped later. both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.
policy	policy	

status

Name	Type	Description
code	string	Encryption progress message code.
message	string	Encryption progress message.

encryption

Name	Type	Description
enabled	boolean	Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate.
key_id	string	The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key.
rekey	boolean	If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	

Name	Type	Description
type	string	Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption). aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption).

error_state

Name	Type	Description
has_bad_blocks	boolean	Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned.
is_inconsistent	boolean	Indicates whether the file system has any inconsistencies. true ‐ File system is inconsistent. false ‐ File system is not inconsistent.

files

Name	Type	Description
maximum	integer	The maximum number of files (inodes) for user-visible data allowed on the volume. This value can be increased or decreased. Increasing the maximum number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and this value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume. Valid in PATCH.
used	integer	Number of files (inodes) used for user-visible data permitted on the volume. This field is valid only when the volume is online.

guarantee

Name	Type	Description
honored	boolean	Is the space guarantee of this volume honored in the aggregate?
type	string	The type of space guarantee of this volume in the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.

Name	Type	Description
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

flexcache

Performance number for FlexCache used to measure cache effectiveness.

Name	Type	Description
cache_miss_percent	integer	Cache miss percentage.
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS, latency and throughput.

Name	Type	Description
_links	_links	
cloud	cloud	Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
flexcache	flexcache	Performance number for FlexCache used to measure cache effectiveness.
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

destination_aggregate

Aggregate

Name	Type	Description
_links	_links	
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
destination_aggregate	destination_aggregate	Aggregate
percent_complete	integer	Completion percentage
state	string	State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress.
tiering_policy	string	Tiering policy for FabricPool

export_policy

Export Policy

Name	Type	Description
_links	_links	
id	integer	
name	string	

nas

Name	Type	Description
export_policy	export_policy	Export Policy

Name	Type	Description
gid	integer	The UNIX group ID of the volume. Valid in POST or PATCH.
path	string	The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline removes its junction path. This attribute is reported in GET only when the volume is mounted.
security_style	string	Security style associated with the volume. Valid in POST or PATCH. mixed ‐ Mixed-style security ntfs ‐ NTFS/Windows-style security unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions unix ‐ Unix-style security.
uid	integer	The UNIX user ID of the volume. Valid in POST or PATCH.

Name	Type	Description
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero.

policy

When "min_throughput_iops", "min_throughput_mbps", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

Name	Type	Description
_links	_links	
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

qos

QoS information

Name	Type	Description
policy	policy	When "min_throughput_iops", "min_throughput_mbps", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

Name	Type	Description
enabled	boolean	This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off".
state	string	Quota state of the volume

retention

Name	Type	Description
default	string	<p>Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num></p>

Name	Type	Description
maximum	string	<p>Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</p>

Name	Type	Description
minimum	string	Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num>

snaplock

Name	Type	Description
append_mode_enabled	boolean	Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires.

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.

Name	Type	Description
privileged_delete	string	Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled.
retention	retention	
type	string	The SnapLock type of the volume. compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files. enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume. non_snaplock ‐ Indicates the volume is non-snaplock.
unspecified_retention_file_count	integer	Indicates the number of files with an unspecified retention time in the volume.

snapmirror

Specifies attributes for SnapMirror protection.

Name	Type	Description
is_protected	boolean	Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data.

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.
enforcement	boolean	Specifies whether space accounting for operations on the volume is done along with storage efficiency.
reporting	boolean	Specifies whether space reporting on the volume is done along with storage efficiency.
used_by_afs	integer	The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.

snapshot

Name	Type	Description
autodelete_enabled	boolean	Specifies whether Snapshot copy autodelete is currently enabled on this volume.
reserve_percent	integer	The space that has been set aside as a reserve for Snapshot copy usage, in percent.
used	integer	The total space used by Snapshot copies in the volume, in bytes.

space

Name	Type	Description
available	integer	The available space, in bytes.
block_storage_inactive_user_data	integer	The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store.
capacity_tier_footprint	integer	Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.
footprint	integer	Data used for this volume in the aggregate, in bytes.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
metadata	integer	Space used by the volume metadata in the aggregate, in bytes.
over_provisioned	integer	The amount of space not available for this volume in the aggregate, in bytes.
performance_tier_footprint	integer	Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.
size	integer	Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.
snapshot	snapshot	
total_footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
used	integer	The virtual space used (includes volume reserves) before storage efficiency, in bytes.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

Name	Type	Description
cache_miss_blocks	integer	Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage.
client_requested_blocks	integer	Total blocks requested by the client.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
cloud	cloud	These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.
flexcache_raw	flexcache_raw	Performance numbers for FlexCache used to measure cache effectiveness.
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

Name	Type	Description
min_cooling_days	integer	<p>This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data.</p> <p>This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.</p>
object_tags	array[string]	<p>This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value".</p>

Name	Type	Description
policy	string	Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks. auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store none ‐ Volume blocks will not be tiered to the cloud store. snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update volume attributes

PATCH /storage/volumes/{uuid}

Introduced In: 9.6

Updates the attributes of a volume. For movement, use the "validate_only" field on the request to validate but not perform the operation. The PATCH API can be used to enable or disable quotas for a FlexVol or a FlexGroup volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline removes its junction path.

Related ONTAP commands

- volume unmount
- volume mount
- volume online
- volume offline
- volume modify
- volume clone modify
- volume efficiency modify
- volume quota on
- volume quota off
- volume snaplock modify

Parameters

Name	Type	In	Required	Description
uuid	string	path	True	Unique identifier of the volume.

Name	Type	In	Required	Description
restore_to.snapshot.uuid	string	query	False	UUID of the Snapshot copy to restore volume to the point in time the Snapshot copy was taken.
restore_to.snapshot.name	string	query	False	Name of the Snapshot copy to restore volume to the point in time the Snapshot copy was taken.
sizing_method	string	query	False	<p>Represents the method to modify the size of a Flexgroup. The following methods are supported:</p> <ul style="list-style-type: none"> • use_existing_resources - Increases or decreases the size of the FlexGroup by increasing or decreasing the size of the current FlexGroup resources • add_new_resources - Increases the size of the FlexGroup by adding new resources • Default value: 1 • enum: ["use_existing_resources", "add_new_resources"]

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
validate_only	boolean	query	False	<p>Validate the operation and its parameters, without actually performing the operation.</p>

Request Body

Name	Type	Description
access_time_enabled	boolean	Indicates whether or not access time updates are enabled on the volume.
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.

Name	Type	Description
analytics	analytics	
application	application	
autosize	autosize	
clone	clone	
cloud_retrieval_policy	string	<p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are "default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.</p> <p>"on_read" policy retrieves tiered data for all client driven data reads.</p> <p>"never" policy never retrieves tiered data.</p> <p>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p>
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents_per_aggregate	integer	Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup. If a volume is being created on a single aggregate, the system will create a flexible volume if the "constituents_per_aggregate" field is not specified, and a FlexGroup if it is specified. If a volume is being created on multiple aggregates, the system will always create a FlexGroup.
create_time	string	Creation time of the volume. This field is generated when the volume is created.
efficiency	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.

Name	Type	Description
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapmirror	snapmirror	Specifies attributes for SnapMirror protection.
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. Taking a volume offline removes its junction path. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6

Example request

```
{
  "aggregates": [
    {
      "name": "aggr1",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    }
  ],
  "analytics": {
    "scan_progress": "17",
    "state": "string",
    "unsupported_reason": {
      "code": "111411207",
      "message": "File system analytics cannot be enabled on volumes
that contain LUNs."
    }
  },
  "application": {
    "name": "string",
    "uuid": "1cd8a442-86d1-11e0-ae1d-123478563412"
  },
  "autosize": {
    "mode": "string"
  },
  "clone": {
    "parent_snapshot": {
      "name": "this_snapshot",
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
    },
    "parent_svm": {
      "name": "svm1",
      "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
    },
    "parent_volume": {
      "name": "volume1",
      "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    },
    "split_complete_percent": 0,
    "split_estimate": 0
  },
  "cloud_retrieval_policy": "string",
  "comment": "string",
  "consistency_group": {
    "name": "consistency_group_1"
  },
}
```

```

"create_time": "2018-06-04T19:00:00Z",
"efficiency": {
  "compaction": "string",
  "compression": "string",
  "cross_volume_dedupe": "string",
  "dedupe": "string",
  "policy": {
    "name": "string"
  }
},
"encryption": {
  "key_id": "string",
  "state": "string",
  "status": {
    "code": "string",
    "message": "string"
  },
  "type": "string"
},
"files": {
  "used": 0
},
"flexcache_endpoint_type": "string",
"guarantee": {
  "type": "string"
},
"movement": {
  "cutover_window": "30",
  "destination_aggregate": {
    "name": "aggr1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "percent_complete": 0,
  "state": "replicating",
  "tiering_policy": "string"
},
"name": "vol_cs_dept",
"nas": {
  "export_policy": {
    "id": "100",
    "name": "default"
  },
  "path": "/user/my_volume",
  "security_style": "string",
  "unix_permissions": "0755"
},

```

```

"qos": {
  "policy": {
    "max_throughput_iops": "10000",
    "max_throughput_mbps": "500",
    "min_throughput_iops": "2000",
    "min_throughput_mbps": "500",
    "name": "performance",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  }
},
"quota": {
  "state": "string"
},
"snaplock": {
  "append_mode_enabled": "",
  "autocommit_period": "P30M",
  "compliance_clock_time": "2018-06-04T19:00:00Z",
  "expiry_time": "Wed Sep  5 11:02:42 GMT 2018",
  "is_audit_log": 1,
  "litigation_count": "10",
  "privileged_delete": "enabled",
  "retention": {
    "default": "P30Y",
    "maximum": "P30Y",
    "minimum": "P30Y"
  },
  "type": "enterprise",
  "unspecified_retention_file_count": "10"
},
"snapshot_policy": {
  "name": "default",
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"space": {
  "available": 0,
  "block_storage_inactive_user_data": 0,
  "capacity_tier_footprint": 0,
  "footprint": 0,
  "local_tier_footprint": 0,
  "logical_space": {
    "available": 0,
    "used_by_afs": 0
  },
  "metadata": 0,
  "over_provisioned": 0,
  "performance_tier_footprint": 0,

```

```

    "snapshot": {
      "used": 0
    },
    "total_footprint": 0,
    "used": 0
  },
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "tiering": {
    "object_tags": [
      "string"
    ],
    "policy": "string",
    "supported": null
  },
  "use_mirrored_aggregates": null,
  "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}

```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```

{
  "job": {
    "uuid": "string"
  }
}

```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
787141	The specified "aggregates.name" and "aggregates.uuid" refer to different aggregates.
917829	Volume autosize grow threshold must be larger than autosize shrink threshold.
917831	Volume minimum autosize must be smaller than the maximum autosize.
918193	Cannot modify tiering min cooling days when vol move is in progress.
918194	Tiering min cooling days not supported for SVMDR.
918195	Tiering min cooling days not supported for non data volumes.
918196	Tiering min cooling days not allowed for the provided tiering policy.
918248	Specifying a value is not valid for initiating volume FlexClone split operation.
918251	Specifying a value is not valid for a Snapshot copy restore operation.
918252	specified "nas.path" is invalid.
918265	Volume is on the same aggregate.
918266	"movement.destination_aggregate" and "movement.state" are mutually exclusive, unless the state is "cutover-wait".
918267	The specified "movement.destination_aggregate" does not exist.
918291	Invalid volume cloud retrieval policy for the provided tiering policy.
918292	cloud retrieval policy not supported for non data volume.
918293	Cannot modify cloud retrieval policy when vol move is in progress.
918521	The volume maximum autosize must be smaller than or equal to the maximum volume size.
13107404	When adding new resources to a FlexGroup by specifying "aggregates.name" or "aggregates.uuid", the FlexGroup cannot be resized using "size". These operations must be done separately.

Error Code	Description
13109187	When adding new resources to a FlexGroup using "sizing_method", "size" must be specified. Neither "aggregates.name" nor "aggregates.uuid" are allowed to be specified, as the aggregates are selected automatically by the system.
13109198	Resizing by adding new resources is only supported for FlexGroups.
111411201	File system analytics cannot be enabled on the target volume because of the specified reason.
111411202	File system analytics cannot be disabled on the target volume because of the specified reason.
111411205	File system analytics requires an effective cluster version of 9.8 or later.
111411206	The specified "analytics.state" is invalid.
111411207	File system analytics cannot be enabled on volumes that contain LUNs.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

aggregates

Aggregate

Name	Type	Description
name	string	
uuid	string	

unsupported_reason

Name	Type	Description
code	string	If file system analytics is not supported on the volume, this field provides the error code explaining why.
message	string	If file system analytics is not supported on the volume, this field provides the error message explaining why.

analytics

Name	Type	Description
scan_progress	integer	Percentage of files in the volume that the file system analytics initialization scan has processed. Only returned when the state is <code>initializing</code> .

Name	Type	Description
state	string	<p>File system analytics state of the volume. If this value is "on", ONTAP collects extra file system analytics information for all directories on the volume. There will be a slight impact to I/O performance to collect this information. If this value is "off", file system analytics information is not collected and not available to be viewed. If this value is "initializing", that means file system analytics was recently turned on, and the initialization scan to gather information all existing files and directories is currently running. If this value is 'unknown' that means there was an internal error when determining the file system analytics state for the volume.</p> <ul style="list-style-type: none"> enum: ["unknown", "initializing", "off", "on"] Introduced in: 9.8
supported	boolean	This field indicates whether or not file system analytics is supported on the volume. If file system analytics is not supported, the reason will be specified in the "analytics.unsupported_reason" field.
unsupported_reason	unsupported_reason	

application

Name	Type	Description
name	string	Name of the application to which the volume belongs. Available only when the volume is part of an application.
uuid	string	UUID of the application to which the volume belongs. Available only when the volume is part of an application.

autosize

Name	Type	Description
grow_threshold	integer	Used space threshold size, in percentage, for the automatic growth of the volume. When the amount of used space in the volume becomes greater than this threshold, the volume automatically grows unless it has reached the maximum size. The volume grows when 'space.used' is greater than this percent of 'space.size'. The 'grow_threshold' size cannot be less than or equal to the 'shrink_threshold' size..
maximum	integer	Maximum size in bytes up to which a volume grows automatically. This size cannot be less than the current volume size, or less than or equal to the minimum size of volume.
minimum	integer	Minimum size in bytes up to which the volume shrinks automatically. This size cannot be greater than or equal to the maximum size of volume.
mode	string	Autosize mode for the volume. grow ‐ Volume automatically grows when the amount of used space is above the 'grow_threshold' value. grow_shrink ‐ Volume grows or shrinks in response to the amount of space used. off ‐ Autosizing of the volume is disabled.

Name	Type	Description
shrink_threshold	integer	Used space threshold size, in percentage, for the automatic shrinkage of the volume. When the amount of used space in the volume drops below this threshold, the volume automatically shrinks unless it has reached the minimum size. The volume shrinks when the 'space.used' is less than the 'shrink_threshold' percent of 'space.size'. The 'shrink_threshold' size cannot be greater than or equal to the 'grow_threshold' size.

snapshot_reference

Name	Type	Description
name	string	
uuid	string	

parent_svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

parent_volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> example: 028baa66-41bd-11e9-81d5-00a0986138f7 Introduced in: 9.6

clone

Name	Type	Description
is_flexclone	boolean	Specifies if this volume is a normal FlexVol or FlexClone. This field needs to be set when creating a FlexClone. Valid in POST.
parent_snapshot	snapshot_reference	
parent_svm	parent_svm	
parent_volume	parent_volume	
split_complete_percent	integer	Percentage of FlexClone blocks split from its parent volume.
split_estimate	integer	Space required by the containing-aggregate to split the FlexClone volume.
split_initiated	boolean	This field is set when split is executed on any FlexClone, that is when the FlexClone volume is split from its parent FlexVol. This field needs to be set for splitting a FlexClone from FlexVol. Valid in PATCH.

consistency_group

Consistency group the volume is part of.

Name	Type	Description
name	string	Name of the consistency group.

policy

Name	Type	Description
name	string	Specifies the name of the efficiency policy.

efficiency

Name	Type	Description
compaction	string	The system can be enabled/disabled compaction. inline ‐ Data will be compacted first and written to the volume. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compaction enabled and some are disabled.
compression	string	The system can be enabled/disabled compression. inline ‐ Data will be compressed first and written to the volume. background ‐ Data will be written to the volume and compressed later. both ‐ Inline compression compresses the data and write to the volume, background compression compresses only the blocks on which inline compression is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are compression enabled and some are disabled.

Name	Type	Description
cross_volume_dedupe	string	The system can be enabled/disabled cross volume dedupe. it can be enabled only when dedupe is enabled. inline ‐ Data will be cross volume deduped first and written to the volume. background ‐ Data will be written to the volume and cross volume deduped later. both ‐ Inline cross volume dedupe dedupes the data and write to the volume, background cross volume dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are cross volume dedupe enabled and some are disabled.
dedupe	string	The system can be enabled/disabled dedupe. inline ‐ Data will be deduped first and written to the volume. background ‐ Data will be written to the volume and deduped later. both ‐ Inline dedupe dedupes the data and write to the volume, background dedupe dedupes only the blocks on which inline dedupe is not run. none ‐ None mixed ‐ Read only field for FlexGroups, where some of the constituent volumes are dedupe enabled and some are disabled.
policy	policy	

status

Name	Type	Description
code	string	Encryption progress message code.
message	string	Encryption progress message.

encryption

Name	Type	Description
enabled	boolean	Creates an encrypted or an unencrypted volume. For POST, when set to 'true', a new key is generated and used to encrypt the given volume. In that case, the underlying SVM must be configured with the key manager. When set to 'false', the volume created will be unencrypted. For PATCH, when set to 'true', it encrypts an unencrypted volume. Specifying the parameter as 'false' in a PATCH operation for an encrypted volume is only supported when moving the volume to another aggregate.
key_id	string	The key ID used for creating encrypted volume. A new key-id is generated for creating an encrypted volume. This key-id is associated with the generated key.
rekey	boolean	If set to 'true', re-encrypts the volume with a new key. Valid in PATCH.
state	string	Volume encryption state. encrypted ‐ The volume is completely encrypted. encrypting ‐ Encryption operation is in progress. partial ‐ Some constituents are encrypted and some are not. Applicable only for FlexGroup volume. rekeying. Encryption of volume with a new key is in progress. unencrypted ‐ The volume is a plain-text one.
status	status	

Name	Type	Description
type	string	Volume encryption type. none ‐ The volume is a plain-text one. volume ‐ The volume is encrypted with NVE (NetApp Volume Encryption). aggregate ‐ The volume is encrypted with NAE (NetApp Aggregate Encryption).

error_state

Name	Type	Description
has_bad_blocks	boolean	Indicates whether the volume has any corrupt data blocks. If the damaged data block is accessed, an IO error, such as EIO for NFS or STATUS_FILE_CORRUPT for CIFS, is returned.
is_inconsistent	boolean	Indicates whether the file system has any inconsistencies. true ‐ File system is inconsistent. false ‐ File system is not inconsistent.

files

Name	Type	Description
maximum	integer	The maximum number of files (inodes) for user-visible data allowed on the volume. This value can be increased or decreased. Increasing the maximum number of files does not immediately cause additional disk space to be used to track files. Instead, as more files are created on the volume, the system dynamically increases the number of disk blocks that are used to track files. The space assigned to track files is never freed, and this value cannot be decreased below the current number of files that can be tracked within the assigned space for the volume. Valid in PATCH.
used	integer	Number of files (inodes) used for user-visible data permitted on the volume. This field is valid only when the volume is online.

guarantee

Name	Type	Description
honored	boolean	Is the space guarantee of this volume honored in the aggregate?
type	string	The type of space guarantee of this volume in the aggregate.

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.

Name	Type	Description
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

cloud

Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.

Name	Type	Description
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

flexcache

Performance number for FlexCache used to measure cache effectiveness.

Name	Type	Description
cache_miss_percent	integer	Cache miss percentage.
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

metric

Performance numbers, such as IOPS, latency and throughput.

Name	Type	Description
cloud	cloud	Performance numbers (IOPS and latency) for cloud store. These numbers are relevant only for volumes hosted on FabricPools.
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
flexcache	flexcache	Performance number for FlexCache used to measure cache effectiveness.
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

destination_aggregate

Aggregate

Name	Type	Description
name	string	
uuid	string	

movement

Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.

Name	Type	Description
cutover_window	integer	Time window in seconds for cutover. The allowed range is between 30 to 300 seconds.
destination_aggregate	destination_aggregate	Aggregate
percent_complete	integer	Completion percentage
state	string	State of volume move operation. PATCH the state to "aborted" to abort the move operation. PATCH the state to "cutover" to trigger cutover. PATCH the state to "paused" to pause the volume move operation in progress. PATCH the state to "replicating" to resume the paused volume move operation. PATCH the state to "cutover_wait" to go into cutover manually. When volume move operation is waiting to go into "cutover" state, this is indicated by the "cutover_pending" state. A change of state is only supported if volume movement is in progress.
tiering_policy	string	Tiering policy for FabricPool

export_policy

Export Policy

Name	Type	Description
id	integer	
name	string	

nas

Name	Type	Description
export_policy	export_policy	Export Policy
gid	integer	The UNIX group ID of the volume. Valid in POST or PATCH.

Name	Type	Description
path	string	The fully-qualified path in the owning SVM's namespace at which the volume is mounted. The path is case insensitive and must be unique within a SVM's namespace. Path must begin with '/' and must not end with '/'. Only one volume can be mounted at any given junction path. An empty path in POST creates an unmounted volume. An empty path in PATCH deactivates and unmounts the volume. Taking a volume offline removes its junction path. This attribute is reported in GET only when the volume is mounted.
security_style	string	Security style associated with the volume. Valid in POST or PATCH. mixed ‐ Mixed-style security ntfs ‐ NTFS/Windows-style security unified ‐ Unified-style security, unified UNIX, NFS and CIFS permissions unix ‐ Unix-style security.
uid	integer	The UNIX user ID of the volume. Valid in POST or PATCH.

Name	Type	Description
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write) and 1 (execute). First digit selects the set user ID(4), set group ID (2) and sticky (1) attributes. The second digit selects permission for the owner of the file; the third selects permissions for other users in the same group; the fourth for other users not in the group. Valid in POST or PATCH. For security style "mixed" or "unix", the default setting is 0755 in octal (493 in decimal) and for security style "ntfs", the default setting is 0000. In cases where only owner, group and other permissions are given (as in 755, representing the second, third and fourth digit), first digit is assumed to be zero.

policy

When "min_throughput_iops", "min_throughput_mbps", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

Name	Type	Description
max_throughput_iops	integer	Specifies the maximum throughput in IOPS, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
max_throughput_mbps	integer	Specifies the maximum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.

Name	Type	Description
min_throughput_iops	integer	Specifies the minimum throughput in IOPS, 0 means none. Setting "min_throughput" is supported on AFF platforms only, unless FabricPool tiering policies are set. This is mutually exclusive with name and UUID during POST and PATCH.
min_throughput_mbps	integer	Specifies the minimum throughput in Megabytes per sec, 0 means none. This is mutually exclusive with name and UUID during POST and PATCH.
name	string	The QoS policy group name. This is mutually exclusive with UUID and other QoS attributes during POST and PATCH.
uuid	string	The QoS policy group UUID. This is mutually exclusive with name and other QoS attributes during POST and PATCH.

qos

QoS information

Name	Type	Description
policy	policy	When "min_throughput_iops", "min_throughput_mbps", "max_throughput_iops" or "max_throughput_mbps" attributes are specified, the storage object is assigned to an auto-generated QoS policy group. If the attributes are later modified, the auto-generated QoS policy-group attributes are modified. Attributes can be removed by specifying "0" and policy group by specifying "none". Upon deletion of the storage object or if the attributes are removed, then the QoS policy-group is also removed.

quota

Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.

Name	Type	Description
enabled	boolean	This option is used to enable or disable the quota for the volume. This option is valid only in PATCH. Quotas are enabled for FlexVols or FlexGroup volumes when the quota state is "on". Quotas are disabled for FlexVols or FlexGroup volumes when the quota state is "off".
state	string	Quota state of the volume

retention

Name	Type	Description
default	string	<p>Specifies the default retention period that is applied to files while committing them to the WORM state without an associated retention period. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period and the string "unspecified" to set an unspecified retention period.</num></num></num></num></num></p>

Name	Type	Description
maximum	string	Specifies the maximum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.

Name	Type	Description
minimum	string	Specifies the minimum allowed retention period for files committed to the WORM state on the volume. The retention value represents a duration and must be specified in the ISO-8601 duration format. The retention period can be in years, months, days, hours, and minutes. A duration specified for years, month,s and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The retention string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the duration field also accepts the string "infinite" to set an infinite retention period.</num></num></num></num></num>

snaplock

Name	Type	Description
append_mode_enabled	boolean	Specifies if the volume append mode is enabled or disabled. When it is enabled, all the files created with write permissions on the volume are, by default, WORM appendable files. The user can append the data to a WORM appendable file but cannot modify the existing contents of the file nor delete the file until it expires.

Name	Type	Description
autocommit_period	string	Specifies the autocommit period for SnapLock volume. All files which are not modified for a period greater than the autocommit period of the volume are committed to the WORM state. The autocommit period value represents a duration and must be specified in the ISO-8601 duration format. The autocommit period can be in years, months, days, hours, and minutes. A period specified for years, months, and days is represented in the ISO-8601 format as "P<num>Y", "P<num>M", "P<num>D" respectively, for example "P10Y" represents a duration of 10 years. A duration in hours and minutes is represented by "PT<num>H" and "PT<num>M" respectively. The period string must contain only a single time element that is, either years, months, days, hours, or minutes. A duration which combines different periods is not supported, for example "P1Y10M" is not supported. Apart from the duration specified in the ISO-8601 format, the autocommit field also accepts the string "none".
compliance_clock_time	string	This is the volume compliance clock time which is used to manage the SnapLock objects in the volume.
expiry_time	string	Expiry time of the volume.
is_audit_log	boolean	Indicates if this volume has been configured as SnapLock audit log volume for the SVM .
litigation_count	integer	Litigation count indicates the number of active legal-holds on the volume.

Name	Type	Description
privileged_delete	string	Specifies the privileged-delete attribute of a SnapLock volume. On a SnapLock Enterprise (SLE) volume, a designated privileged user can selectively delete files irrespective of the retention time of the file. SLE volumes can have privileged delete as disabled, enabled or permanently_disabled and for SnapLock Compliance (SLC) volumes it is always permanently_disabled.
retention	retention	
type	string	The SnapLock type of the volume. compliance ‐ A SnapLock Compliance(SLC) volume provides the highest level of WORM protection and an administrator cannot destroy a SLC volume if it contains unexpired WORM files. enterprise ‐ An administrator can delete a SnapLock Enterprise(SLE) volume. non_snaplock ‐ Indicates the volume is non-snaplock.
unspecified_retention_file_count	integer	Indicates the number of files with an unspecified retention time in the volume.

snapmirror

Specifies attributes for SnapMirror protection.

Name	Type	Description
is_protected	boolean	Specifies whether a volume is a SnapMirror source volume, using SnapMirror to protect its data.

snapshot_policy

This is a reference to the Snapshot copy policy.

Name	Type	Description
name	string	
uuid	string	

logical_space

Name	Type	Description
available	integer	The amount of space available in this volume with storage efficiency space considered used, in bytes.
enforcement	boolean	Specifies whether space accounting for operations on the volume is done along with storage efficiency.
reporting	boolean	Specifies whether space reporting on the volume is done along with storage efficiency.
used_by_afs	integer	The virtual space used by AFS alone (includes volume reserves) and along with storage efficiency, in bytes.

snapshot

Name	Type	Description
autodelete_enabled	boolean	Specifies whether Snapshot copy autodelete is currently enabled on this volume.
reserve_percent	integer	The space that has been set aside as a reserve for Snapshot copy usage, in percent.
used	integer	The total space used by Snapshot copies in the volume, in bytes.

space

Name	Type	Description
available	integer	The available space, in bytes.

Name	Type	Description
block_storage_inactive_user_data	integer	The size that is physically used in the block storage of the volume and has a cold temperature. In bytes. This parameter is only supported if the volume is in an aggregate that is either attached to a cloud store or could be attached to a cloud store.
capacity_tier_footprint	integer	Space used by capacity tier for this volume in the FabricPool aggregate, in bytes.
footprint	integer	Data used for this volume in the aggregate, in bytes.
local_tier_footprint	integer	Space used by the local tier for this volume in the aggregate, in bytes.
logical_space	logical_space	
metadata	integer	Space used by the volume metadata in the aggregate, in bytes.
over_provisioned	integer	The amount of space not available for this volume in the aggregate, in bytes.
performance_tier_footprint	integer	Space used by the performance tier for this volume in the FabricPool aggregate, in bytes.
size	integer	Total provisioned size. The default size is equal to the minimum size of 20MB, in bytes.
snapshot	snapshot	
total_footprint	integer	Data and metadata used for this volume in the aggregate, in bytes.
used	integer	The virtual space used (includes volume reserves) before storage efficiency, in bytes.

iops_raw

The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency_raw

The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

cloud

These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.

Name	Type	Description
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

flexcache_raw

Performance numbers for FlexCache used to measure cache effectiveness.

Name	Type	Description
cache_miss_blocks	integer	Blocks retrieved from origin in case of a cache miss. This can be divided by the raw client_requested_blocks and multiplied by 100 to calculate the cache miss percentage.
client_requested_blocks	integer	Total blocks requested by the client.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
timestamp	string	The timestamp of the performance data.

throughput_raw

Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

statistics

These are raw performance numbers, such as IOPS latency and throughput. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster.

Name	Type	Description
cloud	cloud	These are raw performance numbers (IOPS and latency) for the cloud store. These numbers are aggregated across all nodes in the cluster and increase with the uptime of the cluster. These numbers are relevant only for volumes hosted on FabricPools.
flexcache_raw	flexcache_raw	Performance numbers for FlexCache used to measure cache effectiveness.
iops_raw	iops_raw	The number of I/O operations observed at the storage object. This can be used along with delta time to calculate the rate of I/O operations per unit of time.
latency_raw	latency_raw	The raw latency in microseconds observed at the storage object. This can be divided by the raw IOPS value to calculate the average latency per I/O operation.

Name	Type	Description
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput_raw	throughput_raw	Throughput bytes observed at the storage object. This can be used along with delta time to calculate the rate of throughput bytes per unit of time.
timestamp	string	The timestamp of the performance data.

svm

SVM containing the volume. Required on POST.

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

tiering

Name	Type	Description
min_cooling_days	integer	<p>This parameter specifies the minimum number of days that user data blocks of the volume must be cooled before they can be considered cold and tiered out to the cloud tier. Note that this parameter is only used for tiering purposes and does not affect the reporting of inactive data. The value specified should be greater than the frequency with which applications in the volume shift between different sets of data.</p> <p>This parameter cannot be set when volume tiering policy is either "none" or "all". The default value of this parameter depends on the volume's tiering policy. See the tiering policy section of this documentation for corresponding default values. If the tiering policy on the volume gets changed, then this parameter will be reset to the default value corresponding to the new tiering policy.</p>
object_tags	array[string]	<p>This parameter specifies tags of a volume for objects stored on a FabricPool-enabled aggregate. Each tag is a key,value pair and should be in the format "key=value".</p>

Name	Type	Description
policy	string	Policy that determines whether the user data blocks of a volume in a FabricPool will be tiered to the cloud store when they become cold. FabricPool combines flash (performance tier) with a cloud store into a single aggregate. Temperature of a volume block increases if it is accessed frequently and decreases when it is not. Valid in POST or PATCH. all ‐ This policy allows tiering of both Snapshot copies and active file system user data to the cloud store as soon as possible by ignoring the temperature on the volume blocks. auto ‐ This policy allows tiering of both snapshot and active file system user data to the cloud store none ‐ Volume blocks will not be tiered to the cloud store. snapshot_only ‐ This policy allows tiering of only the volume Snapshot copies not associated with the active file system. The default tiering policy is "snapshot-only" for a FlexVol and "none" for a FlexGroup. The default minimum cooling period for the "snapshot-only" tiering policy is 2 days and for the "auto" tiering policy is 31 days.

volume

Name	Type	Description
access_time_enabled	boolean	Indicates whether or not access time updates are enabled on the volume.
aggregates	array[aggregates]	Aggregate hosting the volume. Required on POST.
analytics	analytics	
application	application	
autosize	autosize	

Name	Type	Description
clone	clone	
cloud_retrieval_policy	string	<p>This parameter specifies the cloud retrieval policy for the volume. This policy determines which tiered out blocks to retrieve from the capacity tier to the performance tier. The available cloud retrieval policies are</p> <p>"default" policy retrieves tiered data based on the underlying tiering policy. If the tiering policy is 'auto', tiered data is retrieved only for random client driven data reads. If the tiering policy is 'none' or 'snapshot_only', tiered data is retrieved for random and sequential client driven data reads. If the tiering policy is 'all', tiered data is not retrieved.</p> <p>"on_read" policy retrieves tiered data for all client driven data reads.</p> <p>"never" policy never retrieves tiered data.</p> <p>"promote" policy retrieves all eligible tiered data automatically during the next scheduled scan. It is only supported when the tiering policy is 'none' or 'snapshot_only'. If the tiering policy is 'snapshot_only', the only data brought back is the data in the AFS. Data that is only in a snapshot copy stays in the cloud and if tiering policy is 'none' then all data is retrieved.</p>
comment	string	A comment for the volume. Valid in POST or PATCH.
consistency_group	consistency_group	Consistency group the volume is part of.

Name	Type	Description
constituents_per_aggregate	integer	Specifies the number of times to iterate over the aggregates listed with the "aggregates.name" or "aggregates.uuid" when creating or expanding a FlexGroup. If a volume is being created on a single aggregate, the system will create a flexible volume if the "constituents_per_aggregate" field is not specified, and a FlexGroup if it is specified. If a volume is being created on multiple aggregates, the system will always create a FlexGroup.
create_time	string	Creation time of the volume. This field is generated when the volume is created.
efficiency	efficiency	
encryption	encryption	
error_state	error_state	
files	files	
flexcache_endpoint_type	string	FlexCache endpoint type. none ‐ The volume is neither a FlexCache nor origin of any FlexCache. cache ‐ The volume is a FlexCache volume. origin ‐ The volume is origin of a FlexCache volume.
guarantee	guarantee	
is_object_store	boolean	Specifies whether the volume is provisioned for an object store server.
is_svm_root	boolean	Specifies whether the volume is a root volume of the SVM it belongs to.

Name	Type	Description
movement	movement	Volume movement. All attributes are modify, that is, not writable through POST. Set PATCH state to destination_aggregate to initiate a volume move operation. Volume movement on FlexGroup constituents are not supported.
name	string	Volume name. The name of volume must start with an alphabetic character (a to z or A to Z) or an underscore (_). The name must be 197 or fewer characters in length for FlexGroups, and 203 or fewer characters in length for all other types of volumes. Volume names must be unique within an SVM. Required on POST.
nas	nas	
qos	qos	QoS information
quota	quota	Quotas track the space or file usage of a user, group, or qtree in a FlexVol or a FlexGroup volume.
size	integer	Physical size of the volume, in bytes. The minimum size for a FlexVol volume is 20MB and the minimum size for a FlexGroup volume is 200MB per constituent. The recommended size for a FlexGroup volume is a minimum of 100GB per constituent. For all volumes, the default size is equal to the minimum size.
snaplock	snaplock	
snapmirror	snapmirror	Specifies attributes for SnapMirror protection.
snapshot_policy	snapshot_policy	This is a reference to the Snapshot copy policy.
space	space	

Name	Type	Description
state	string	Volume state. A volume can only be brought online if it is offline. Taking a volume offline removes its junction path. The 'mixed' state applies to FlexGroup volumes only and cannot be specified as a target state. An 'error' state implies that the volume is not in a state to serve data.
svm	svm	SVM containing the volume. Required on POST.
tiering	tiering	
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • readOnly: 1 • Introduced in: 9.6

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage files and directories

Storage volumes volume.uuid files path endpoint overview

Overview

This API is used to read a file, write to a file, retrieve a list of files and directories, and retrieve or modify certain properties of files and directories. The path field is used to specify the path to the directory or file to be acted on. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.

File data

Read and write data from/to a named file. To read a file, the Accept request HTTP header must be specified as multipart/form-data, and a value for the length query property, which represents the number of bytes to be read, must be specified. The API will fail if the length of data being read/written exceeds 1 MB. This API should only be used on normal files or streams associated with files. The results for other file types, such as LUNs is undefined.

The following APIs are used to read or write data to a file:

– GET /api/storage/volumes/{volume.uuid}/files/{path}?byte_offset=0&length=40 -H "Accept: multipart/form-data"

– POST /api/storage/volumes/{volume.uuid}/files/{path} -H "Content-Type: multipart/form-data" --form "file=the data to be written to the new file"

– PATCH /api/storage/volumes/{volume.uuid}/files/{path}?byte_offset=10 -H "Content-Type: multipart/form-data" --form "file=the new data to be written or overwritten to the existing file starting at byte_offset"

Listing directories and files

A list of files and directories and their properties can be retrieved for a specified path.

The following APIs are used to view a list of files and directories:

– GET /api/storage/volumes/{volume.uuid}/files

– GET /api/storage/volumes/{volume.uuid}/files/{path}

– GET /api/storage/volumes/{volume.uuid}/files/{path}?fields=*

File information

The metadata and detailed information about a single directory or file can be retrieved by setting the `return_metadata` query property to `true`. The information returned includes `type`, `creation_time`, `modified_time`, `changed_time`, `accessed_time`, `unix_permissions`, `owner_id`, `group_id`, `size`, `hard_links_count`, `inode_number`, `is_empty`, `bytes_used`, `unique_bytes`, `inode_generation`, `is_vm_aligned`, `is_junction`, `links`, and `analytics` (if requested).

The following API is used to view the properties of a single file or directory:

– GET `/api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true`

File usage

Custom details about the usage of a file can be retrieved by specifying a value for the `byte_offset` and `length` query properties.

The following API is used to view the unique bytes, and bytes used, by a file based on the range defined by `byte_offset` and `length`:

– GET
`/api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true&byte_offset={int}&length={int}`

Create a directory

The following API is used to create a directory:

– POST `/api/storage/volumes/{volume.uuid}/files/{path} -d { "type" : "directory", "unix-permissions" : "644" }`

Delete an entire directory

A directory can be deleted. The behavior of this call is equivalent to `rm -rf`.

The following API is used to delete an entire directory:

– DELETE `/api/storage/volumes/{volume.uuid}/files/{path}?recurse=true`

Delete a file or an empty directory

The following API is used to delete a file or an empty directory:

– DELETE `/api/storage/volumes/{volume.uuid}/files/{path}`

– DELETE `/api/storage/volumes/{volume.uuid}/files/{path}?recurse=false`

File system analytics

File system analytics provide a quick method for obtaining information summarizing properties of all files within any directory tree of a volume. When file system analytics are enabled on a volume, `analytics.*` fields may be requested, and will be populated in the response records corresponding to directories. The API does not support file system analytics for requests that are made beyond the boundary of the specified `volume.uuid`.

The following APIs are used to obtain analytics information for a directory:

– GET `/api/storage/volumes/{volume.uuid}/files/{path}?fields=analytics`

– GET /api/storage/volumes/{volume.uuid}/files/{path}?fields=**

QoS

QoS policies and settings enforce Service Level Objectives (SLO) on a file. A pre-created QoS policy can be used by specifying the `qos.name` or `qos.uuid` properties.

The following APIs are used to assign a QoS policy to a file:

– PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "qos_policy.name" : "policy" }'

– PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "qos_policy.uuid" : "b89bc5dd-94a3-11e8-a7a3-0050568edf84" }'

Symlinks

The following APIs are used to create a symlink and read the contents of a symlink:

– POST /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "target" : "directory2/file1" }'

– GET /api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true&fields=target

Rename a file or a directory

The following API can be used to rename a file or a directory. Note that you need to provide the path relative to the root of the volume in the `path` body parameter.

– PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "path" : "directory1/directory2" }'

– PATCH /api/storage/volumes/{volume.uuid}/files/{path} -d '{ "path" : "directory1/directory2/file1" }'

Examples

Writing to a new file

```
# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-90f9-005056a73aff/files/aNewFile" -H "Content-Type: multipart/form-data" --form "file=the data to be written to the new file"
```

Writing to an existing file

```
# The API:
PATCH /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-90f9-005056a73aff/files/aNewFile?byte_offset=39" -H "Content-Type: multipart/form-data" --form "file=*here is a little more data"
```

Reading a file

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-90f9-005056a73aff/files/aNewFile?byte_offset=0&length=100" -H "Accept: multipart/form-data"

# Response for file data:
--ec51b3541741ade7
Content-Disposition: form-data; name="bytes_read"
Content-Type: text/plain
66
--ec51b3541741ade7
Content-Disposition: form-data; filename="aNewFile"
Content-Type: application/octet-stream
the data to be written to the new file*here is a little more data
--ec51b3541741ade7--
```

Creating a directory

You can use the POST request to create a directory.

```
# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/dir1" -H 'accept: application/hal+json' -d '{
  "type" : "directory", "unix_permissions" : "644" }'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "path": "dir1",
      "type": "directory",
      "unix_permissions": 644
    }
  ]
}
```

Creating a stream on a file

```
# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}?overwrite=true

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/54c06ce2-5430-11ea-90f9-005056a73aff/files/aNewFile?overwrite=true&byte_offset=-1&stream_name=someStream" -H "Content-Type: multipart/form-data" --form "file=the data to be written to the new file"
```

Retrieving the list of files in a directory

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/d1%2Fd2%2Fd3"

# Response for file records:
{
  "records": [
```

```

{
  "path": "d1/d2/d3",
  "name": ".",
  "type": "directory",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2F%2E"
    },
    "metadata": {
      "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E?return_metadata=true"
    }
  }
},
{
  "path": "d1/d2/d3",
  "name": "..",
  "type": "directory",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2F%2E%2E"
    },
    "metadata": {
      "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E%2E?return_metadata=true"
    }
  }
},
{
  "path": "d1/d2/d3",
  "name": "f1",
  "type": "file",
  "_links": {
    "metadata": {
      "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2Ffile1?return_metadata=true"
    }
  }
},
{
  "path": "d1/d2/d3",
  "name": "d5",
  "type": "directory",
  "_links": {

```

```

    "self": {
      "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2Fd5"
    },
    "metadata": {
      "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2Fd5?return_metadata=true"
    }
  }
},
"num_records": 4,
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3"
  }
}
}

```

Retrieving a list of files based on file type

You can filter the list of files you retrieve based on multiple file types by including a query parameter in the following format `type="file|symlink"`

```

# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/d1%2Fd2%2Fd3?type=file&#124;directory"

# Response for file records:
{
  "records": [
    {
      "path": "d1/d2/d3",
      "name": ".",
      "type": "directory",
      "_links": {
        "self": {
          "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2F%2E"
        },
        "metadata": {
          "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-

```

```

005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E?return_metadata=true"
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "..",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2F%2E%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2F%2E%2E?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "f1",
    "type": "file",
    "_links": {
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2Ffile1?return_metadata=true"
      }
    }
  },
  {
    "path": "d1/d2/d3",
    "name": "d5",
    "type": "directory",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d1%2Fd2%2Fd3%2Fd5"
      },
      "metadata": {
        "href": "/api/storage/volumes/e8274d79-3bba-11ea-b780-005056a7d72a/files/d1%2Fd2%2Fd3%2Fd5?return_metadata=true"
      }
    }
  }
],

```

```
"num_records": 4,  
"_links": {  
  "self": {  
    "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-  
005056aca658/files/d1%2Fd2%2Fd3"  
  }  
}  
}
```

Retrieving the properties of a directory or a file

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}?return_metadata=true

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/d1%2Fd2%2Fd3%2Ff1?return_metadata=true"

# Response for file properties:
{
  "records": [
    {
      "path": "d1/d2/d3/f1",
      "name": "",
      "type": "file",
      "creation_time": "2019-06-12T21:27:28-04:00",
      "modified_time": "2019-06-12T21:27:28-04:00",
      "changed_time": "2019-06-12T21:27:28-04:00",
      "accessed_time": "2019-06-12T21:27:28-04:00",
      "unix_permissions": 644,
      "owner_id": 54738,
      "group_id": 30,
      "size": 200,
      "hard_links_count": 1,
      "inode_number": 1233,
      "bytes_used": 4096,
      "unique_bytes": 4096,
      "inode_generation": 214488325,
      "is_vm_aligned": false,
      "is_junction": false
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/storage/volumes/da8bb06c-823e-11e9-b790-005056acdcb0/files/d1%2Fd2%2Fd3%2Ff1?return_metadata=true"
    }
  }
}
```

Creating a symlink to a relative path

You can use the POST request to create a symlink.


```
# The API:
POST /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X POST "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/symlink1" -H 'accept: application/hal+json' -d '{
"target" : "d1/f1"}'

# The response:
{
  "num_records": 1,
  "records": [
    {
      "path": "symlink1",
      "target": "d1/f1"
    }
  ]
}
```

Retrieving the target of a symlink

You can use the GET request to view the target of a symlink.

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/symlink1?return_metadata=true&fields=target"

# The response:
{
  "records": [
    {
      "path": "symlink1",
      "target": "dl/fl"
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/storage/volumes/54c06ce2-5430-11ea-90f9-
005056a73aff/files/symlink1?return_metadata=true&fields=target"
    }
  }
}
```

Retrieving the usage information for a file

You can use the GET request to retrieve the unique bytes held in a file with or without specifying the offset.

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/f1?return_metadata=true&byte_offset=100&length=200"

# The response:
{
  "records": [
    {
      "path": "d1/d2/d3/f1",
      "type": "file",
      "creation_time": "2019-06-12T21:27:28-04:00",
      "modified_time": "2019-06-12T21:27:28-04:00",
      "changed_time": "2019-06-12T21:27:28-04:00",
      "accessed_time": "2019-06-12T21:27:28-04:00",
      "unix_permissions": 644,
      "owner_id": 54738,
      "group_id": 30,
      "size": 200,
      "hard_links_count": 1,
      "inode_number": 1233,
      "bytes_used": 4096,
      "unique_bytes": 4096,
      "inode_generation": 214488325,
      "is_vm_aligned": false,
      "is_junction": false
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/storage/volumes/cb6b139-8d21-11e9-b926-05056aca658/files/f1?return_metadata=true&byte_offset=100&length=200"
    }
  }
}
```

Retrieving all information (including analytics) for a directory

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860/files/d1?return_metadata=true&fields=**"

# Response for all fields of the directory:
{
  "records": [
    {
      "svm": {
        "uuid": "58a996a2-f9d5-11e9-8043-00505682f860",
        "_links": {
          "self": {
            "href": "/api/svm/svms/58a996a2-f9d5-11e9-8043-00505682f860"
          }
        }
      },
      "volume": {
        "uuid": "1ef5d1b2-f9d7-11e9-8043-00505682f860",
        "_links": {
          "self": {
            "href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860"
          }
        }
      },
      "path": "d1",
      "type": "directory",
      "creation_time": "2019-10-28T23:04:13+00:00",
      "modified_time": "2019-10-28T23:10:30+00:00",
      "changed_time": "2019-10-28T23:10:30+00:00",
      "accessed_time": "2019-10-28T23:10:38+00:00",
      "unix_permissions": 755,
      "owner_id": 1002,
      "group_id": 65533,
      "size": 4096,
      "hard_links_count": 5,
      "inode_number": 96,
      "is_empty": false,
      "bytes_used": 4096,
      "inode_generation": 214514951,
      "is_vm_aligned": false,
      "is_junction": false,
      "analytics": {
        "file_count": 668,
        "bytes_used": 209657856,
        "subdir_count": 18,

```

```

"by_modified_time": {
  "bytes_used": {
    "values": [
      0,
      0,
      0,
      0,
      3112960,
      0,
      14041088,
      20545536,
      0,
      57933824,
      61947904,
      68804608,
      188686336,
      0,
      0,
      0,
      20971520,
      0
    ],
    "percentages": [
      0,
      0,
      0,
      0,
      1.48,
      0,
      6.7,
      9.8,
      0,
      27.63,
      29.55,
      32.82,
      90,
      0,
      0,
      0,
      10,
      0
    ],
    "labels": [
      "2019-W42",
      "2019-W41",
      "2019-W40",

```

```

        "2019-W39",
        "2019-W38",
        "2019-10",
        "2019-09",
        "2019-08",
        "2019-Q4",
        "2019-Q3",
        "2019-Q2",
        "2019-Q1",
        "2019",
        "2018",
        "2017",
        "2016",
        "--2015",
        "unknown"
    ]
}
},
"by_accessed_time": {
    "bytes_used": {
        "values": [
            102760448,
            1867776,
            1245184,
            2179072,
            1556480,
            105873408,
            9027584,
            8093696,
            105873408,
            23969792,
            32382976,
            26460160,
            188686336,
            0,
            0,
            0,
            20971520,
            0
        ],
        "percentages": [
            49.01,
            0.89,
            0.59,
            1.04,
            0.74,

```

```

        50.5,
        4.31,
        3.86,
        50.5,
        11.43,
        15.45,
        12.62,
        90,
        0,
        0,
        0,
        10,
        0
    ],
    "labels": [
        "2019-W42",
        "2019-W41",
        "2019-W40",
        "2019-W39",
        "2019-W38",
        "2019-10",
        "2019-09",
        "2019-08",
        "2019-Q4",
        "2019-Q3",
        "2019-Q2",
        "2019-Q1",
        "2019",
        "2018",
        "2017",
        "2016",
        "--2015",
        "unknown"
    ]
}
}
}
},
"num_records": 1,
"_links": {
    "self": {
        "href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860/files/d1?return_metadata=true&fields=**"
    }
}
}

```

```
}
```

Retrieving file system analytics information for a set of histogram buckets

```
# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-
05056aca658/files/d3?type=directory&fields=analytics&analytics.histogram_b
y_time_labels=2019-Q3,2019-Q2,2019-Q1,2018-Q4"

# Response with analytics data
{
  "records": [
    {
      "path": "d3",
      "name": ".",
      "type": "directory",
      "analytics": {
        "file_count": 44,
        "bytes_used": 244240384,
        "subdir_count": 14,
        "by_modified_time": {
          "bytes_used": {
            "values": [
              57344,
              29720576,
              196141056,
              57344
            ],
            "percentages": [
              0.02,
              12.17,
              80.31,
              0.02
            ]
          }
        },
        "by_accessed_time": {
          "bytes_used": {
            "values": [
              69632,
              244170752,
```



```

        0,
        0
    ],
    "percentages": [
        0.03,
        99.97,
        0,
        0
    ]
}
}
},
"_links": {
    "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2F%2E"
    },
    "metadata": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2F%2E?return_metadata=true"
    }
}
},
{
    "path": "d3",
    "name": "..",
    "type": "directory",
    "analytics": {
        "file_count": 515,
        "bytes_used": 3034574848,
        "subdir_count": 23,
        "by_modified_time": {
            "bytes_used": {
                "values": [
                    61440,
                    1756479488,
                    214622208,
                    1191936
                ],
                "percentages": [
                    0,
                    57.88,
                    7.07,
                    0.04
                ]
            }
        }
    }
}

```

```

    },
    "by_accessed_time": {
      "bytes_used": {
        "values": [
          282624,
          3034292224,
          0,
          0
        ],
        "percentages": [
          0.01,
          99.99,
          0,
          0
        ]
      }
    },
    "_links": {
      "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2F%2E%2E"
      },
      "metadata": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2F%2E%2E?return_metadata=true"
      }
    },
    {
      "path": "d3",
      "name": "d5",
      "type": "directory",
      "analytics": {
        "file_count": 10,
        "bytes_used": 47648768,
        "subdir_count": 4,
        "by_modified_time": {
          "bytes_used": {
            "values": [
              0,
              29638656,
              0,
              0
            ],
            "percentages": [

```

```

        0,
        62.20,
        0,
        0
    ]
}
},
"by_accessed_time": {
    "bytes_used": {
        "values": [
            0,
            47648768,
            0,
            0
        ],
        "percentages": [
            0,
            100,
            0,
            0
        ]
    }
}
},
"_links": {
    "self": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2Fd5"
    },
    "metadata": {
        "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3%2Fd5?return_metadata=true"
    }
}
},
"num_records": 3,
"analytics": {
    "by_modified_time": {
        "bytes_used": {
            "labels": [
                "2019-Q3",
                "2019-Q2",
                "2019-Q1",
                "2018-Q4"
            ]
        }
    }
}
}

```

```

    }
  },
  "by_accessed_time": {
    "bytes_used": {
      "labels": [
        "2019-Q3",
        "2019-Q2",
        "2019-Q1",
        "2018-Q4"
      ]
    }
  }
},
"_links": {
  "self": {
    "href": "/api/storage/volumes/cb6b1b39-8d21-11e9-b926-005056aca658/files/d3?type=directory&fields=analytics&analytics.histogram_by_time_labels=2019-Q3,2019-Q2,2019-Q1,2018-Q4"
  }
}
}

```

Identifying the largest subdirectories

```

# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860/files/d1?fields=analytics.bytes_used&type=directory&order_by=analytics.bytes_used%20desc"

# Response with the largest subdirectories sorted by their usage:
{
  "records": [
    {
      "path": "d1",
      "name": "..",
      "type": "directory",
      "analytics": {
        "bytes_used": 56623104
      }
    },
    {

```

```

    "path": "d1",
    "name": ".",
    "type": "directory",
    "analytics": {
      "bytes_used": 35651584
    }
  },
  {
    "path": "d1",
    "name": "biggest",
    "type": "directory",
    "analytics": {
      "bytes_used": 17825792
    }
  },
  {
    "path": "d1",
    "name": "bigger",
    "type": "directory",
    "analytics": {
      "bytes_used": 10485760
    }
  },
  {
    "path": "d1",
    "name": "big",
    "type": "directory",
    "analytics": {
      "bytes_used": 5242880
    }
  }
],
"num_records": 5,
"_links": {
  "self": {
    "href": "/api/storage/volumes/1ef5d1b2-f9d7-11e9-8043-00505682f860/files/d1?fields=analytics.bytes_used&type=directory&order_by=analytics.bytes_used%20desc"
  }
}
}

```

Assigning a QoS policy to a file

You can use the PATCH request to assign a QoS policy to a file.

```
# The API:
PATCH /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-
b926-05056aca658/files/directory1%2Ffile1" -d '{ "qos_policy": { "name":
"policy" } }'

# The response:
{ }
```

Retrieving QoS information for a file

You can use the GET request for all fields with `return_metadata="true"` to retrieve QoS information for the file.

```

# The API:
GET /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/file?return_metadata=true&fields=**"

# The response:
{
  "records": [
    {
      "svm": {
        "uuid": "42ee3002-67dd-11ea-8508-005056a7b8ac"
      },
      "volume": {
        "uuid": "c05eb66a-685f-11ea-8508-005056a7b8ac"
      },
      "path": "file",
      "type": "lun",
      "creation_time": "2020-03-17T10:58:40-04:00",
      "modified_time": "2020-03-24T18:15:40-04:00",
      "changed_time": "2020-03-24T18:15:40-04:00",
      "accessed_time": "2020-03-24T18:15:40-04:00",
      "unix_permissions": 644,
      "owner_id": 0,
      "group_id": 0,
      "size": 1048576,
      "hard_links_count": 2,
      "inode_number": 96,
      "bytes_used": 1056768,
      "inode_generation": 219748425,
      "is_vm_aligned": false,
      "is_junction": false,
      "is_snapshot": false,
      "qos_policy": {
        "name": "pg1",
        "uuid": "00725264-688f-11ea-8f10-005056a7b8ac"
      }
    }
  ],
  "num_records": 1
}

```

Deleting an entire directory

You can use the DELETE request to remove an entire directory recursively.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Fdirectory2?recurse=true"

# The response:
{
  "job": {
    "uuid": "27d287e8-fcd4-11e9-b8a4-005056a7b97b",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/27d287e8-fcd4-11e9-b8a4-005056a7b97b"
      }
    }
  }
}
```

Deleting an entire directory with specified throttling threshold

You can specify the maximum number of directory delete operations per second when removing an entire directory recursively.


```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Fdirectory2?recurse=true&throttle-deletion=100"

# The response:
{
  "job": {
    "uuid": "27d287e8-fcd4-11e9-b8a4-005056a7b97b",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/27d287e8-fcd4-11e9-b8a4-005056a7b97b"
      }
    }
  }
}
```

Deleting an empty directory

You can use the DELETE request to remove an empty directory.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Fdirectory2"

# The response:
{}
```

Deleting a file

You can use the DELETE request to remove a file.

```
# The API:
DELETE /api/storage/volumes/{volume.uuid}/files/{path}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/cb6b1b39-8d21-11e9-b926-05056aca658/files/directory1%2Ffile2"

# The response:
{ }
```

Delete an existing file or directory

DELETE /storage/volumes/{volume.uuid}/files/{path}

Introduced In: 9.8

Deletes an existing file or directory. Query-based DELETE operations are not supported.

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
path	string	path	True	The relative path of a directory in the volume. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.
recurse	boolean	query	False	Delete an entire directory. The behaviour of this call is equivalent to <code>rm -rf</code> . <ul style="list-style-type: none"> • Default value:
throttle-deletion	integer	query	False	The maximum number of directory delete operations per second. A valid throttle-deletion number is an interger from 10 to 100000.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 200, Ok

Response

Status: 202, Accepted

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
131074	No such file or directory.
131102	Read-only file system.
131138	Directory not empty.
918235	A volume with UUID {volume.uuid} was not found.
6488081	The {field} field is not supported for DELETE operations.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve files and directories

GET /storage/volumes/{volume.uuid}/files/{path}

Introduced In: 9.7

Retrieves a list of files and directories for a given directory or returns only the properties of a single given directory or file of a volume.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query property. See [Requesting specific fields](#) to learn more.

- `analytics`
- `qos_policy.name`
- `qos_policy.uuid`

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID

Name	Type	In	Required	Description
path	string	path	True	Relative path of a file or directory in the volume. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.
byte_offset	integer	query	False	The file offset to start reading from. <ul style="list-style-type: none"> Introduced in: 9.8
length	integer	query	False	Length of the range in bytes. <ul style="list-style-type: none"> Introduced in: 9.8
return_metadata	boolean	query	False	If true, the request returns metadata for the the directory or file specified in the path. <ul style="list-style-type: none"> Introduced in: 9.8 Default value:
fill_enabled	boolean	query	False	Filter by fill_enabled <ul style="list-style-type: none"> Introduced in: 9.8
bytes_used	integer	query	False	Filter by bytes_used
is_snapshot	boolean	query	False	Filter by is_snapshot <ul style="list-style-type: none"> Introduced in: 9.8
owner_id	integer	query	False	Filter by owner_id

Name	Type	In	Required	Description
overwrite_enabled	boolean	query	False	Filter by overwrite_enabled <ul style="list-style-type: none"> • Introduced in: 9.8
unique_bytes	integer	query	False	Filter by unique_bytes <ul style="list-style-type: none"> • Introduced in: 9.8
changed_time	string	query	False	Filter by changed_time
type	string	query	False	Filter by type
is_junction	boolean	query	False	Filter by is_junction
path	string	query	False	Filter by path
size	integer	query	False	Filter by size
is_vm_aligned	boolean	query	False	Filter by is_vm_aligned
inode_number	integer	query	False	Filter by inode_number
analytics.by_accessed_time.bytes_used.oldest_label	string	query	False	Filter by analytics.by_accessed_time.bytes_used.oldest_label <ul style="list-style-type: none"> • Introduced in: 9.8
analytics.by_accessed_time.bytes_used.values	integer	query	False	Filter by analytics.by_accessed_time.bytes_used.values <ul style="list-style-type: none"> • Introduced in: 9.8

Name	Type	In	Required	Description
analytics.by_accessed_time.bytes_used.newest_label	string	query	False	Filter by analytics.by_accessed_time.bytes_used.newest_label • Introduced in: 9.8
analytics.by_accessed_time.bytes_used.labels	string	query	False	Filter by analytics.by_accessed_time.bytes_used.labels • Introduced in: 9.8
analytics.by_accessed_time.bytes_used.percentages	number	query	False	Filter by analytics.by_accessed_time.bytes_used.percentages • Introduced in: 9.8
analytics.file_count	integer	query	False	Filter by analytics.file_count • Introduced in: 9.8
analytics.bytes_used	integer	query	False	Filter by analytics.bytes_used • Introduced in: 9.8
analytics.by_modified_time.bytes_used.oldest_label	string	query	False	Filter by analytics.by_modified_time.bytes_used.oldest_label • Introduced in: 9.8

Name	Type	In	Required	Description
analytics.by_modified_time.bytes_used.values	integer	query	False	Filter by analytics.by_modified_time.bytes_used.values • Introduced in: 9.8
analytics.by_modified_time.bytes_used.newest_label	string	query	False	Filter by analytics.by_modified_time.bytes_used.newest_label • Introduced in: 9.8
analytics.by_modified_time.bytes_used.labels	string	query	False	Filter by analytics.by_modified_time.bytes_used.labels • Introduced in: 9.8
analytics.by_modified_time.bytes_used.percentages	number	query	False	Filter by analytics.by_modified_time.bytes_used.percentages • Introduced in: 9.8
analytics.subdir_count	integer	query	False	Filter by analytics.subdir_count • Introduced in: 9.8
name	string	query	False	Filter by name
target	string	query	False	Filter by target • Introduced in: 9.8
accessed_time	string	query	False	Filter by accessed_time

Name	Type	In	Required	Description
qos_policy.name	string	query	False	Filter by qos_policy.name • Introduced in: 9.8
qos_policy.uuid	string	query	False	Filter by qos_policy.uuid • Introduced in: 9.8
modified_time	string	query	False	Filter by modified_time
inode_generation	integer	query	False	Filter by inode_generation
unix_permissions	integer	query	False	Filter by unix_permissions
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
hard_links_count	integer	query	False	Filter by hard_links_count
group_id	integer	query	False	Filter by group_id
is_empty	boolean	query	False	Filter by is_empty
creation_time	string	query	False	Filter by creation_time

Name	Type	In	Required	Description
analytics.histogram_by_time_labels	array[string]	query	False	<p>Request that returned analytics_histogram_by_time objects including values associated with the specified labels.</p> <p>As described in the object description, the labels may take the following forms:partial-date
<tt>--</tt>partial-date
partial-date <tt>--</tt>
partial-date<tt>.-</tt>partial-date
<tt>unknown</tt>Intervals that the system would not normally return may be specified. In this case, the appropriate values and percentages summarizing all files with a time-based attribute within the indicated period of time are calculated and returned in the response. However, there are some restrictions:Any partial-date specified as the beginning or end of an interval must be tracked by the system. Valid partial-dates may be determined by making an OPTIONS request to</p>

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
analytics	analytics	<p>Additional file system analytics information that is invariant amongst all elements in the collection.</p> <p>This property is only populated if file system analytics is enabled on the containing volume.</p> <p>This analytics object captures properties that are invariant amongst all elements included in the <code>records</code> array. The invariant properties are included here, rather than within the information for each element, to avoid returning an excessive amount of duplicated information when the collection is large.</p>
num_records	integer	Number of records.
records	array[file_info]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ]
      }
    },
    "by_modified_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ]
      }
    }
  },
  "records": [
    {
      "_links": {
        "metadata": {
          "href": "/api/resourcelink"
        },
        "self": {
```

```

    "href": "/api/resourcelink"
  }
},
"accessed_time": "2019-06-12T11:00:16-04:00",
"analytics": {
  "by_accessed_time": {
    "bytes_used": {
      "labels": [
        "2019-07",
        "2019-06",
        "2019-05",
        "2019",
        "2018",
        "--2017",
        "unknown"
      ],
      "newest_label": [
        "2019-07",
        "2019-06",
        "2019-05",
        "2019",
        "2018",
        "--2017",
        "unknown"
      ],
      "oldest_label": [
        "2019-07",
        "2019-06",
        "2019-05",
        "2019",
        "2018",
        "--2017",
        "unknown"
      ],
      "percentages": [
        "0.1",
        "11.24",
        "0.18",
        "15.75",
        "0.75",
        "83.5",
        "0"
      ],
      "values": [
        "15925248",
        "1735569408",

```

```

        "27672576",
        "2430595072",
        "116105216",
        "12889948160",
        "0"
    ]
}
},
"by_modified_time": {
    "bytes_used": {
        "labels": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "newest_label": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "oldest_label": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "percentages": [
            "0.1",
            "11.24",
            "0.18",
            "15.75",
            "0.75",
            "83.5",
            "0"
        ],
    },

```



```

        "values": [
            "15925248",
            "1735569408",
            "27672576",
            "2430595072",
            "116105216",
            "12889948160",
            "0"
        ]
    },
    "bytes_used": "15436648448",
    "file_count": "21134",
    "subdir_count": "35"
},
"bytes_used": "4096",
"changed_time": "2019-06-12T11:00:16-04:00",
"creation_time": "2019-06-12T11:00:16-04:00",
"group_id": "30",
"hard_links_count": "1",
"inode_generation": "214753547",
"inode_number": "1695",
"is_empty": "",
"is_junction": "",
"is_snapshot": "",
"is_vm_aligned": "",
"modified_time": "2019-06-12T11:00:16-04:00",
"name": "test_file",
"owner_id": "54738",
"path": "d1/d2/d3",
"qos_policy": {
    "_links": {
        "self": {
            "href": "/api/resourcelink"
        }
    },
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "200",
"target": "some_directory/some_other_directory/some_file",
"type": "file",
"unique_bytes": "4096",
"unix_permissions": "0755",
"volume": {
    "_links": {

```

```

        "self": {
            "href": "/api/resourcelink"
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
    }
}
]
}

```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```

{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}

```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

bytes_used

Number of bytes used on-disk, broken down by date of last access.



Name	Type	Description
labels	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

by_accessed_time

File system analytics information, broken down by date of last access.


interval includes points in time arbitrarily far in the past. When

Name	Type	Description
bytes_used	bytes_used	Number of bytes used on-disk, broken down by date of last access.

bytes_used

with any other time period. This usually occurs when the data was at some point associated with a time in the future.

Number of bytes used on-disk, broken down by date of last modification.



Name	Type	Description
labels	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

by_modified_time

File system analytics information, broken down by date of last modification. This interval includes points in time arbitrarily far in the past. When

Name	Type	Description
bytes_used	bytes_used	Number of bytes used on-disk, broken down by date of last modification.

analytics

with any other time period. This usually occurs when the data was at some point associated with a

Additional file system analytics information that is invariant amongst all elements in the collection.

This property is only populated if file system analytics is enabled on the containing volume.

This analytics object captures properties that are invariant amongst all elements included in the `records` array. The invariant properties are included here, rather than within the information for each element, to avoid returning an excessive amount of duplicated information when the collection is large.

Name	Type	Description
by_accessed_time	by_accessed_time	File system analytics information, broken down by date of last access.
by_modified_time	by_modified_time	File system analytics information, broken down by date of last modification.

_links

Name	Type	Description
metadata	href	
self	href	

bytes_used

Number of bytes used on-disk, broken down by date of last access.



Name	Type	Description
labels	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
newest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>


Name	Type	Description
oldest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
percentages	array[number]	Percentages for this histogram
values	array[integer]	Values for this histogram

bytes_used

Number of bytes used on-disk, broken down by date of last modification

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.



Name	Type	Description
labels	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59. <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59. Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
newest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
oldest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
percentages	array[number]	Percentages for this histogram
values	array[integer]	Values for this histogram

analytics

Additional file system analytics information summarizing all descendants of a directory with any other time period.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a [GET /storage/volumes/{volume.uuid}/files/{path}](#) call returns a large collection.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was at some point associated with a time in the future.

Name	Type	Description
by_accessed_time	by_accessed_time	File system analytics information, broken down by date of last access.
by_modified_time	by_modified_time	File system analytics information, broken down by date of last modification.
bytes_used	integer	Number of bytes used on-disk
file_count	integer	Number of descendants
subdir_count	integer	Number of sub directories

_links

Name	Type	Description
self	href	

qos_policy

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".



Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on `/storage/luns` to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

Name	Type	Description
_links	_links	
name	string	The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string "" or set it to "none" in a PATCH request.
uuid	string	The unique identifier of the QoS policy. Valid in PATCH.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	<p>Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move.</p> <ul style="list-style-type: none"> • example: 028baa66-41bd-11e9-81d5-00a0986138f7 • Introduced in: 9.6


file_info

Information about a single file.

Name	Type	Description
_links	_links	
accessed_time	string	Last access time of the file in date-time format.

Name	Type	Description
analytics	analytics	<p>Additional file system analytics information summarizing all descendents of a directory.</p> <p>This property is only populated if file system analytics is enabled on the containing volume.</p> <p>In the context of the <code>records</code> property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection.</p>
bytes_used	integer	The actual number of bytes used on disk by this file. If <code>byte_offset</code> and <code>length</code> parameters are specified, this will return the bytes used by the file within the given range.
changed_time	string	Last time data or attributes changed on the file in date-time format.
creation_time	string	Creation time of the file in date-time format.
fill_enabled	boolean	Returns "true" if the space reservation is enabled. The field <code>overwrite_enabled</code> must also be set to the same value as this field.
group_id	integer	The integer ID of the group of the file owner.
hard_links_count	integer	The number of hard links to the file.
inode_generation	integer	Inode generation number.
inode_number	integer	The file inode number.

Name	Type	Description
is_empty	boolean	Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.
is_junction	boolean	Returns "true" if the directory is a junction.
is_snapshot	boolean	Returns "true" if the directory is a Snapshot copy.
is_vm_aligned	boolean	Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.
modified_time	string	Last data modification time of the file in date-time format.
name	string	Name of the file.
overwrite_enabled	boolean	Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field.
owner_id	integer	The integer ID of the file owner.
path	string	Path of the file.

Name	Type	Description
qos_policy	qos_policy	<p>The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both <code>qos_policy.uuid</code> and <code>qos_policy.name</code> properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property <code>qos_policy.name</code> in a PATCH request to an empty string "" or "none".</p> <div>  <p>Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on <code>/storage/luns</code> to assign a QoS policy for such files.</p> </div> <p>Note that a QoS policy can be set on a file, or a file's volume, but not on both.</p>
size	integer	The size of the file, in bytes.
target	string	The relative or absolute path contained in a symlink, in the form <code><some>/<path>.</path></some></code>
type	string	Type of the file.
unique_bytes	integer	Number of bytes uniquely held by this file. If <code>byte_offset</code> and <code>length</code> parameters are specified, this will return bytes uniquely held by the file within the given range.

Name	Type	Description
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Write to an existing file with the supplied data

```
PATCH /storage/volumes/{volume.uuid}/files/{path}
```

Introduced In: 9.8

Writes to an existing file with the supplied data or modifies the size, name, space reservation information, QoS policy, or hole range information of a file. Query-based PATCH operations are not supported.


Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
path	string	path	True	Relative path of a file in the volume. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.
byte_offset	integer	query	False	How many bytes into the file to begin writing. Use -1 to append (default).
overwrite	boolean	query	False	If false, and the file exists, the write will fail. Default is false.
stream_name	string	query	False	Name of stream associated with the file to write data to.
data	string	formData	False	Data to write to the file.

Request Body

Name	Type	Description
accessed_time	string	Last access time of the file in date-time format.
analytics	analytics	<p>Additional file system analytics information summarizing all descendents of a directory.</p> <p>This property is only populated if file system analytics is enabled on the containing volume.</p> <p>In the context of the <code>records</code> property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection.</p>

Name	Type	Description
bytes_used	integer	The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range.
changed_time	string	Last time data or attributes changed on the file in date-time format.
creation_time	string	Creation time of the file in date-time format.
fill_enabled	boolean	Returns "true" if the space reservation is enabled. The field overwrite_enabled must also be set to the same value as this field.
group_id	integer	The integer ID of the group of the file owner.
hard_links_count	integer	The number of hard links to the file.
inode_generation	integer	Inode generation number.
inode_number	integer	The file inode number.
is_empty	boolean	Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.
is_junction	boolean	Returns "true" if the directory is a junction.

Name	Type	Description
is_vm_aligned	boolean	Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.
modified_time	string	Last data modification time of the file in date-time format.
name	string	Name of the file.
overwrite_enabled	boolean	Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field.
owner_id	integer	The integer ID of the file owner.
qos_policy	qos_policy	<p>The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both qos_policy.uuid and qos_policy.name properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property qos_policy.name in a PATCH request to an empty string "" or "none".</p> <div>  <p>Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.</p> </div> <p>Note that a QoS policy can be set on a file, or a file's volume, but not on both.</p>

Name	Type	Description
size	integer	The size of the file, in bytes.
target	string	The relative or absolute path contained in a symlink, in the form <some>/<path>.</path></some>
unique_bytes	integer	Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range.
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.
volume	volume	

Example request

```
{
  "accessed_time": "2019-06-12T11:00:16-04:00",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "oldest_label": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "percentages": [
          "0.1",
          "11.24",
          "0.18",
          "15.75",
          "0.75",
          "83.5",
          "0"
        ],
        "values": [
          "15925248",
          "1735569408",
```

```

        "27672576",
        "2430595072",
        "116105216",
        "12889948160",
        "0"
    ]
}
},
"by_modified_time": {
    "bytes_used": {
        "labels": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "newest_label": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "oldest_label": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "percentages": [
            "0.1",
            "11.24",
            "0.18",
            "15.75",
            "0.75",
            "83.5",
            "0"
        ],
    },

```

```

        "values": [
            "15925248",
            "1735569408",
            "27672576",
            "2430595072",
            "116105216",
            "12889948160",
            "0"
        ]
    },
    "bytes_used": "15436648448",
    "file_count": "21134",
    "subdir_count": "35"
},
"bytes_used": "4096",
"changed_time": "2019-06-12T11:00:16-04:00",
"creation_time": "2019-06-12T11:00:16-04:00",
"group_id": "30",
"hard_links_count": "1",
"inode_generation": "214753547",
"inode_number": "1695",
"is_empty": "",
"is_junction": "",
"is_snapshot": "",
"is_vm_aligned": "",
"modified_time": "2019-06-12T11:00:16-04:00",
"name": "test_file",
"owner_id": "54738",
"qos_policy": {
    "name": "qos1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
},
"size": "200",
"target": "some_directory/some_other_directory/some_file",
"unique_bytes": "4096",
"unix_permissions": "0755",
"volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}

```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
918235	A volume with UUID {volume.uuid} was not found.
6488081	The {field} field is not supported for PATCH operations.
6488082	Failed to rename {path}.
6488083	Failed to rename {path} to {path} because a directory named {path} already exists.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

bytes_used

Number of bytes used on-disk, broken down by date of last access.



Name	Type	Description
labels	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
newest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
oldest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59. <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59. Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
percentages	array[number]	Percentages for this histogram
values	array[integer]	Values for this histogram

by_accessed_time


File system analytics information, broken down by date of last access.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was

Name	Type	Description
bytes_used	bytes_used	Number of bytes used on-disk, broken down by date of last access.

bytes_used

Number of bytes used on-disk, broken down by date of last modification.



Name	Type	Description
labels	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
newest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
oldest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59. <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
percentages	array[number]	Percentages for this histogram
values	array[integer]	Values for this histogram

by_modified_time

File system analytics information, broken down by date of last modification.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was

Name	Type	Description
bytes_used	bytes_used	Number of bytes used on-disk, broken down by date of last modification.

analytics

Additional file system analytics information summarizing all descendents of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a [GET /storage/volumes/{volume.uuid}/files/{path}](#) call returns a large collection.

Name	Type	Description
by_accessed_time	by_accessed_time	File system analytics information, broken down by date of last access.
by_modified_time	by_modified_time	File system analytics information, broken down by date of last modification.
bytes_used	integer	Number of bytes used on-disk
file_count	integer	Number of descendants
subdir_count	integer	Number of sub directories

qos_policy

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".



Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

Name	Type	Description
name	string	The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string "" or set it to "none" in a PATCH request.
uuid	string	The unique identifier of the QoS policy. Valid in PATCH.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6


file_info

Information about a single file.

Name	Type	Description
accessed_time	string	Last access time of the file in date-time format.

Name	Type	Description
analytics	analytics	<p>Additional file system analytics information summarizing all descendents of a directory.</p> <p>This property is only populated if file system analytics is enabled on the containing volume.</p> <p>In the context of the <code>records</code> property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection.</p>
bytes_used	integer	The actual number of bytes used on disk by this file. If <code>byte_offset</code> and <code>length</code> parameters are specified, this will return the bytes used by the file within the given range.
changed_time	string	Last time data or attributes changed on the file in date-time format.
creation_time	string	Creation time of the file in date-time format.
fill_enabled	boolean	Returns "true" if the space reservation is enabled. The field <code>overwrite_enabled</code> must also be set to the same value as this field.
group_id	integer	The integer ID of the group of the file owner.
hard_links_count	integer	The number of hard links to the file.
inode_generation	integer	Inode generation number.
inode_number	integer	The file inode number.

Name	Type	Description
is_empty	boolean	Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.
is_junction	boolean	Returns "true" if the directory is a junction.
is_vm_aligned	boolean	Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.
modified_time	string	Last data modification time of the file in date-time format.
name	string	Name of the file.
overwrite_enabled	boolean	Returns "true" if the space reservation for overwrites is enabled. The field fill_enabled must also be set to the same value as this field.
owner_id	integer	The integer ID of the file owner.

Name	Type	Description
qos_policy	qos_policy	<p>The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both <code>qos_policy.uuid</code> and <code>qos_policy.name</code> properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property <code>qos_policy.name</code> in a PATCH request to an empty string "" or "none".</p> <div>  <p>Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on <code>/storage/luns</code> to assign a QoS policy for such files.</p> </div> <p>Note that a QoS policy can be set on a file, or a file's volume, but not on both.</p>
size	integer	The size of the file, in bytes.
target	string	<p>The relative or absolute path contained in a symlink, in the form</p> <pre><some>/<path>.</path></some></pre>
unique_bytes	integer	Number of bytes uniquely held by this file. If <code>byte_offset</code> and <code>length</code> parameters are specified, this will return bytes uniquely held by the file within the given range.

Name	Type	Description
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a new file with the supplied data

POST /storage/volumes/{volume.uuid}/files/{path}

Introduced In: 9.8

Creates a new file with the supplied data, creates a new directory or creates a new symlink.

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
path	string	path	True	Relative path of a new file, directory or symlink. The path field requires using "%2E" to represent "." and "%2F" to represent "/" for the path provided.
byte_offset	integer	query	False	How many bytes into the file to begin writing. Use -1 to append (default).
overwrite	boolean	query	False	If false, and the file exists, the write will fail. Default is false.
stream_name	string	query	False	Name of stream associated with the file to write data to.
data	string	formData	False	Data to write to the file.

Request Body

Name	Type	Description
accessed_time	string	Last access time of the file in date-time format.
analytics	analytics	<p>Additional file system analytics information summarizing all descendents of a directory.</p> <p>This property is only populated if file system analytics is enabled on the containing volume.</p> <p>In the context of the <code>records</code> property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection.</p>

Name	Type	Description
bytes_used	integer	The actual number of bytes used on disk by this file. If byte_offset and length parameters are specified, this will return the bytes used by the file within the given range.
changed_time	string	Last time data or attributes changed on the file in date-time format.
creation_time	string	Creation time of the file in date-time format.
group_id	integer	The integer ID of the group of the file owner.
hard_links_count	integer	The number of hard links to the file.
inode_generation	integer	Inode generation number.
inode_number	integer	The file inode number.
is_empty	boolean	Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.
is_junction	boolean	Returns "true" if the directory is a junction.
is_vm_aligned	boolean	Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.

Name	Type	Description
modified_time	string	Last data modification time of the file in date-time format.
name	string	Name of the file.
owner_id	integer	The integer ID of the file owner.
path	string	Path of the file.
target	string	The relative or absolute path contained in a symlink, in the form <some>/<path>.</path></some>
type	string	Type of the file.
unique_bytes	integer	Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range.
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.
volume	volume	

Example request

```
{
  "accessed_time": "2019-06-12T11:00:16-04:00",
  "analytics": {
    "by_accessed_time": {
      "bytes_used": {
        "labels": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "newest_label": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "oldest_label": [
          "2019-07",
          "2019-06",
          "2019-05",
          "2019",
          "2018",
          "--2017",
          "unknown"
        ],
        "percentages": [
          "0.1",
          "11.24",
          "0.18",
          "15.75",
          "0.75",
          "83.5",
          "0"
        ],
        "values": [
          "15925248",
          "1735569408",
```

```

        "27672576",
        "2430595072",
        "116105216",
        "12889948160",
        "0"
    ]
}
},
"by_modified_time": {
    "bytes_used": {
        "labels": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "newest_label": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "oldest_label": [
            "2019-07",
            "2019-06",
            "2019-05",
            "2019",
            "2018",
            "--2017",
            "unknown"
        ],
        "percentages": [
            "0.1",
            "11.24",
            "0.18",
            "15.75",
            "0.75",
            "83.5",
            "0"
        ],
    },

```

```

        "values": [
            "15925248",
            "1735569408",
            "27672576",
            "2430595072",
            "116105216",
            "12889948160",
            "0"
        ]
    },
    "bytes_used": "15436648448",
    "file_count": "21134",
    "subdir_count": "35"
},
"bytes_used": "4096",
"changed_time": "2019-06-12T11:00:16-04:00",
"creation_time": "2019-06-12T11:00:16-04:00",
"fill_enabled": null,
"group_id": "30",
"hard_links_count": "1",
"inode_generation": "214753547",
"inode_number": "1695",
"is_empty": "",
"is_junction": "",
"is_snapshot": "",
"is_vm_aligned": "",
"modified_time": "2019-06-12T11:00:16-04:00",
"name": "test_file",
"overwrite_enabled": null,
"owner_id": "54738",
"path": "d1/d2/d3",
"target": "some_directory/some_other_directory/some_file",
"type": "file",
"unique_bytes": "4096",
"unix_permissions": "0755",
"volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
}
}

```

Response

Status: 201, Created

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
917505	The SVM does not exist.
917525	The volume in the symlink path does not exist in the SVM.
917698	The volume in the symlink path is not mounted in the namespace.
6488064	This command is not supported.
6488065	The volume in the symlink path is invalid.
6488066	Mounting the unjunctioned volume in the symlink path failed.
6488069	Internal file error.
6488084	Failed to create {path} because the "unix_permissions" field was not specified.
6488085	Failed to create {path} because the "type" field was not specified.
8257536	This operation is not supported for the system volume specified in the symlink path.
8257541	Failed to compute the SVM identification from this content.
8257542	This operation is not supported for the administrative SVM.
9437549	This operation is not allowed on SVMs with Infinite Volume.
13172837	This operation is not permitted because the SVM is locked for a migrate operation.

Definitions

See Definitions


href

Name	Type	Description
href	string	

_links

bytes_used

Number of bytes used on-disk, broken down by date of last access.



Name	Type	Description
labels	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59. <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59. Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
newest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
oldest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
percentages	array[number]	Percentages for this histogram
values	array[integer]	Values for this histogram

by_accessed_time

File system analytics information, broken down by date of last access.

- The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was

Name	Type	Description
bytes_used	bytes_used	Number of bytes used on-disk, broken down by date of last access.

bytes_used

Number of bytes used on-disk, broken down by date of last modification.



Name	Type	Description
labels	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
newest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms: a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used: Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
oldest_label	array[string]	<p>Labels for this histogram.</p> <p>Each label is a string indicating the period of time the corresponding data is associated with. Elements of the array take one of the following forms:a partial date in an extended ISO8601 representation an interval between partial dates in an extended ISO8601 representation, where "--" is used to separate the beginning and end of the interval the string literal "unknown"</p> <p>For partial dates and partial date intervals where components of a date are unspecified, the label allows for any valid normalized values the unspecified components might take. For example, the label "2017" allows for any time within the year 2017. Essentially, this is the fully specified interval 2017-01-01T00:00:00—2017-12-31T23:59:59. Similarly, the interval "2018-05—2018-07" allows for any time within the months of May, June, and July in 2018, corresponding to the fully specified interval 2018-05-01T00:00:00—2018-07-31T23:59:59.</p> <p>The following extensions to ISO8601 are used:Quarters may be specified. The form yyyy-Qq is used to represent the qth quarter of the year yyyy. Q1 consists of the months January, February, and March; Q2 consists of April, May, and June; Q3 consists of July, August, and September; Q4 consists of October, November, and December. For example, the label "2019-Q2" represents the second quarter of the year 2019, which corresponds to the interval 2019-04-01T00:00:00—2019-06-30T23:59:59.</p> <p>Either the beginning or end of an interval may be omitted. When the beginning is omitted, the</p>

Name	Type	Description
percentages	array[number]	Percentages for this histogram
values	array[integer]	Values for this histogram

by_modified_time

File system analytics information, broken down by date of last modification.

The "unknown" label tracks data that could not be associated with any other time period. This usually occurs when the data was

Name	Type	Description
bytes_used	bytes_used	Number of bytes used on-disk, broken down by date of last modification.

analytics

Additional file system analytics information summarizing all descendents of a directory.

This property is only populated if file system analytics is enabled on the containing volume.

In the context of the `records` property of a [GET /storage/volumes/{volume.uuid}/files/{path}](#) call returns a large collection.

Name	Type	Description
by_accessed_time	by_accessed_time	File system analytics information, broken down by date of last access.
by_modified_time	by_modified_time	File system analytics information, broken down by date of last modification.
bytes_used	integer	Number of bytes used on-disk
file_count	integer	Number of descendants
subdir_count	integer	Number of sub directories

qos_policy

The QoS policy for the file. Both traditional and adaptive QoS policies are supported. If both `qos_policy.uuid` and `qos_policy.name` properties are specified in the same request, they must refer to the same QoS policy. To remove the file from a QoS policy, set the property `qos_policy.name` in a PATCH request to an empty string "" or "none".



Files which are in use as a LUN cannot be assigned to a QoS policy, instead use PATCH on /storage/luns to assign a QoS policy for such files.

Note that a QoS policy can be set on a file, or a file's volume, but not on both.

Name	Type	Description
name	string	The name of the QoS policy. To remove the file from a QoS policy, set this property to an empty string "" or set it to "none" in a PATCH request.
uuid	string	The unique identifier of the QoS policy. Valid in PATCH.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

file_info

Information about a single file.

Name	Type	Description
accessed_time	string	Last access time of the file in date-time format.

Name	Type	Description
analytics	analytics	<p>Additional file system analytics information summarizing all descendents of a directory.</p> <p>This property is only populated if file system analytics is enabled on the containing volume.</p> <p>In the context of the <code>records</code> property of a GET /storage/volumes/{volume.uuid}/files/{path} call returns a large collection.</p>
bytes_used	integer	The actual number of bytes used on disk by this file. If <code>byte_offset</code> and <code>length</code> parameters are specified, this will return the bytes used by the file within the given range.
changed_time	string	Last time data or attributes changed on the file in date-time format.
creation_time	string	Creation time of the file in date-time format.
group_id	integer	The integer ID of the group of the file owner.
hard_links_count	integer	The number of hard links to the file.
inode_generation	integer	Inode generation number.
inode_number	integer	The file inode number.

Name	Type	Description
is_empty	boolean	Specifies whether or not a directory is empty. A directory is considered empty if it only contains entries for "." and "..". This element is present if the file is a directory. In some special error cases, such as when the volume goes offline or when the directory is moved while retrieving this info, this field might not get set.
is_junction	boolean	Returns "true" if the directory is a junction.
is_vm_aligned	boolean	Returns true if the file is vm-aligned. A vm-aligned file is a file that is initially padded with zero-filled data so that its actual data starts at an offset other than zero. The amount by which the start offset is adjusted depends on the vm-align setting of the hosting volume.
modified_time	string	Last data modification time of the file in date-time format.
name	string	Name of the file.
owner_id	integer	The integer ID of the file owner.
path	string	Path of the file.
target	string	The relative or absolute path contained in a symlink, in the form <some>/<path>.</path></some>
type	string	Type of the file.
unique_bytes	integer	Number of bytes uniquely held by this file. If byte_offset and length parameters are specified, this will return bytes uniquely held by the file within the given range.

Name	Type	Description
unix_permissions	integer	UNIX permissions to be viewed as an octal number. It consists of 4 digits derived by adding up bits 4 (read), 2 (write), and 1 (execute). The first digit selects the set user ID(4), set group ID (2), and sticky (1) attributes. The second digit selects permissions for the owner of the file; the third selects permissions for other users in the same group; the fourth selects permissions for other users not in the group.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve historical performance metrics for a volume

GET /storage/volumes/{volume.uuid}/metrics

Introduced In: 9.7

Retrieves historical performance metrics for a volume.

Parameters

Name	Type	In	Required	Description
latency.total	integer	query	False	Filter by latency.total
latency.read	integer	query	False	Filter by latency.read
latency.other	integer	query	False	Filter by latency.other
latency.write	integer	query	False	Filter by latency.write
cloud.duration	string	query	False	Filter by cloud.duration
cloud.timestamp	string	query	False	Filter by cloud.timestamp
cloud.latency.total	integer	query	False	Filter by cloud.latency.total
cloud.latency.read	integer	query	False	Filter by cloud.latency.read
cloud.latency.other	integer	query	False	Filter by cloud.latency.other
cloud.latency.write	integer	query	False	Filter by cloud.latency.write
cloud.iops.total	integer	query	False	Filter by cloud.iops.total
cloud.iops.read	integer	query	False	Filter by cloud.iops.read
cloud.iops.other	integer	query	False	Filter by cloud.iops.other
cloud.iops.write	integer	query	False	Filter by cloud.iops.write
cloud.status	string	query	False	Filter by cloud.status

Name	Type	In	Required	Description
flexcache.timestamp	string	query	False	Filter by flexcache.timestamp • Introduced in: 9.8
flexcache.cache_miss_percent	integer	query	False	Filter by flexcache.cache_miss_percent • Introduced in: 9.8
flexcache.duration	string	query	False	Filter by flexcache.duration • Introduced in: 9.8
flexcache.status	string	query	False	Filter by flexcache.status • Introduced in: 9.8
status	string	query	False	Filter by status
timestamp	string	query	False	Filter by timestamp
duration	string	query	False	Filter by duration
throughput.total	integer	query	False	Filter by throughput.total
throughput.read	integer	query	False	Filter by throughput.read
throughput.other	integer	query	False	Filter by throughput.other
throughput.write	integer	query	False	Filter by throughput.write
iops.total	integer	query	False	Filter by iops.total
iops.read	integer	query	False	Filter by iops.read

Name	Type	In	Required	Description
iops.other	integer	query	False	Filter by iops.other
iops.write	integer	query	False	Filter by iops.write
volume.uuid	string	path	True	Unique identifier of the volume.
interval	string	query	False	<p>The time range for the data. Examples can be 1h, 1d, 1m, 1w, 1y. The period for each time range is as follows:</p> <ul style="list-style-type: none"> • 1h: Metrics over the most recent hour sampled over 15 seconds. • 1d: Metrics over the most recent day sampled over 5 minutes. • 1w: Metrics over the most recent week sampled over 30 minutes. • 1m: Metrics over the most recent month sampled over 2 hours. • 1y: Metrics over the most recent year sampled over a day. • Default value: 1 • enum: ["1h", "1d", "1w", "1m", "1y"]

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	return_records	boolean	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "duration": "PT15S",
      "iops": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "latency": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "status": "ok",
      "throughput": {
        "read": "200",
        "total": "1000",
        "write": "100"
      },
      "timestamp": "2017-01-25T11:20:13Z"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

iops

The rate of I/O operations observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

latency

The round trip latency in microseconds observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

throughput

The rate of throughput bytes per second observed at the storage object.

Name	Type	Description
other	integer	Performance metric for other I/O operations. Other I/O operations can be metadata operations, such as directory lookups and so on.
read	integer	Performance metric for read I/O operations.
total	integer	Performance metric aggregated over all types of I/O operations.
write	integer	Performance metric for write I/O operations.

records

Performance numbers, such as IOPS latency and throughput.

Name	Type	Description
_links	_links	

Name	Type	Description
duration	string	The duration over which this sample is calculated. The time durations are represented in the ISO-8601 standard format. Samples can be calculated over the following durations:
iops	iops	The rate of I/O operations observed at the storage object.
latency	latency	The round trip latency in microseconds observed at the storage object.
status	string	Errors associated with the sample. For example, if the aggregation of data over multiple nodes fails, then any partial errors might return "ok" on success or "error" on an internal uncategorized failure. Whenever a sample collection is missed but done at a later time, it is back filled to the previous 15 second timestamp and tagged with "backfilled_data". "Inconsistent_delta_time" is encountered when the time between two collections is not the same for all nodes. Therefore, the aggregated value might be over or under inflated. "Negative_delta" is returned when an expected monotonically increasing value has decreased in value. "Inconsistent_old_data" is returned when one or more nodes do not have the latest data.
throughput	throughput	The rate of throughput bytes per second observed at the storage object.
timestamp	string	The timestamp of the performance data.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage volume Snapshot copies

Storage volumes volume.uuid snapshots endpoint overview

Overview

A Snapshot copy is the view of the filesystem as it exists at the time when the Snapshot copy is created.

In ONTAP, different types of Snapshot copies are supported, such as scheduled Snapshot copies, user requested Snapshot copies, SnapMirror Snapshot copies, and so on.

ONTAP Snapshot copy APIs allow you to create, modify, delete and retrieve Snapshot copies.

Snapshot copy APIs

The following APIs are used to perform operations related to Snapshot copies.

– POST /api/storage/volumes/{volume.uuid}/snapshots

– GET /api/storage/volumes/{volume.uuid}/snapshots

– GET /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

– PATCH /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

– DELETE /api/storage/volumes/{volume.uuid}/snapshots/{uuid}

Examples

Creating a Snapshot copy

The POST operation is used to create a Snapshot copy with the specified attributes.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots

# The call:
curl -X POST "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots" -H "accept:
application/hal+json" -d '{"name": "snapshot_copy", "comment": "Store this
copy." }'

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:43:34 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/?name=snapshot_copy
Content-Length: 189
Content-Type: application/json
{
  "num_records": 1,
  "records": [
    {
      "volume": {
        "name": "v2"
      },
      "svm": {
        "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
        "name": "vs0"
      },
      "name": "snapshot_copy",
      "comment": "Store this copy."
    }
  ],
  "job": {
    "uuid": "6f68c85b-45e1-11e9-8fc7-005056bbc848",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6f68c85b-45e1-11e9-8fc7-005056bbc848"
      }
    }
  }
}
```

```

}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:43:57 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 224
Content-Type: application/json
{
  "uuid": "6f68c85b-45e1-11e9-8fc7-005056bbc848",
  "description": "POST /api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/?name=snapshot_copy",
  "state": "success",
  "message": "success",
  "code": 0
}

```

Retrieving Snapshot copy attributes

The GET operation is used to retrieve Snapshot copy attributes.

```

# The API:
/api/storage/volumes/{volume.uuid}/snapshots

# The call:
curl -X GET "https://<mgmt-
ip>/api/storage/volumes/{volume.uuid}/snapshots" -H "accept:
application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 21:14:06 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Type: application/json
Transfer-Encoding: chunked
{
  "records": [
    {
      "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
      "name": "hourly.2019-03-13_1305",
      "_links": {
        "self": {

```

```

      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
    }
  },
  {
    "uuid": "f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8",
    "name": "hourly.2019-03-13_1405",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/f0dd497f-efe8-44b7-a4f4-bdd3890bc0c8"
      }
    }
  },
  {
    "uuid": "02701900-51bd-46b8-9c77-47d9a9e2ce1d",
    "name": "hourly.2019-03-13_1522",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots/02701900-51bd-46b8-9c77-47d9a9e2ce1d"
      }
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848/snapshots"
  }
}
}

```

Retrieving the attributes of a specific Snapshot copy

The GET operation is used to retrieve the attributes of a specific Snapshot copy.

```

# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X GET "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0" -H
"accept: application/hal+json"

# The response:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:39:26 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 308
Content-Type: application/json
{
  "volume": {
    "uuid": "0353dc05-405f-11e9-acb6-005056bbc848",
    "name": "v2",
    "_links": {
      "self": {
        "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-005056bbc848"
      }
    }
  },
  "uuid": "402b6c73-73a0-4e89-a58a-75ee0ab3e8c0",
  "svm": {
    "uuid": "8139f958-3c6e-11e9-a45f-005056bbc848",
    "name": "vs0",
    "_links": {
      "self": {
        "href": "/api/svm/svms/8139f958-3c6e-11e9-a45f-005056bbc848"
      }
    }
  },
  "name": "hourly.2019-03-13_1305",
  "create_time": "2019-03-13T13:05:00-04:00",
  "_links": {
    "self": {
      "href": "/api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/402b6c73-73a0-4e89-a58a-75ee0ab3e8c0"
    }
  }
}

```

Updating a Snapshot copy

The PATCH operation is used to update the specific attributes of a Snapshot copy.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X PATCH "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f" -d
'{"name": "snapshot_copy_new" }' -H "accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:50:44 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
  "job": {
    "uuid": "6f7c3a82-45e2-11e9-8fc7-005056bbc848",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6f7c3a82-45e2-11e9-8fc7-005056bbc848"
      }
    }
  }
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 22:54:16 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 242
Content-Type: application/json
{
  "uuid": "6f7c3a82-45e2-11e9-8fc7-005056bbc848",
  "description": "PATCH /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f",
  "state": "success",
  "message": "success",
  "code": 0
}
```

Deleting a Snapshot copy

The DELETE operation is used to delete a Snapshot copy.

```
# The API:
/api/storage/volumes/{volume.uuid}/snapshots/{uuid}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/storage/volumes/0353dc05-405f-11e9-
acb6-005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f" -H
"accept: application/hal+json"

# The response:
HTTP/1.1 202 Accepted
Date: Wed, 13 Mar 2019 22:57:51 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 189
Content-Type: application/json
{
  "job": {
    "uuid": "6da1dfdd-45e3-11e9-8fc7-005056bbc848",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/6da1dfdd-45e3-11e9-8fc7-005056bbc848"
      }
    }
  }
}

# The Job:
HTTP/1.1 200 OK
Date: Wed, 13 Mar 2019 23:02:46 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 243
Content-Type: application/json
{
  "uuid": "6da1dfdd-45e3-11e9-8fc7-005056bbc848",
  "description": "DELETE /api/storage/volumes/0353dc05-405f-11e9-acb6-
005056bbc848/snapshots/16f7008c-18fd-4a7d-8485-a0e290d9db7f",
  "state": "success",
  "message": "success",
  "code": 0
}
```


Retrieve volume Snapshot copies

GET /storage/volumes/{volume.uuid}/snapshots

Introduced In: 9.6

Retrieves a collection of volume Snapshot copies.

Related ONTAP commands

- `snapshot show`

Learn more

- [DOC /storage/volumes/{volume.uuid}/snapshots](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume
state	string	query	False	Filter by state
owners	string	query	False	Filter by owners <ul style="list-style-type: none">• Introduced in: 9.7
snaplock_expiry_time	string	query	False	Filter by snaplock_expiry_time
svm.uuid	string	query	False	Filter by svm.uuid
svm.name	string	query	False	Filter by svm.name
expiry_time	string	query	False	Filter by expiry_time
volume.uuid	string	query	False	Filter by volume.uuid
volume.name	string	query	False	Filter by volume.name
create_time	string	query	False	Filter by create_time
name	string	query	False	Filter by name
comment	string	query	False	Filter by comment

Name	Type	In	Required	Description
uuid	string	query	False	Filter by uuid
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
num_records	integer	Number of records
records	array[snapshot]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "comment": "string",
      "create_time": "2019-02-04T19:00:00Z",
      "expiry_time": "2019-02-04T19:00:00Z",
      "name": "this_snapshot",
      "owners": [
        "string"
      ],
      "snaplock_expiry_time": "2019-02-04T19:00:00Z",
      "state": "string",
      "svm": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
      "volume": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "volume1",
        "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
      }
    }
  ]
}
```

```
}  
]  
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{  
  "error": {  
    "arguments": [  
      {  
        "code": "string",  
        "message": "string"  
      }  
    ],  
    "code": "4",  
    "message": "entry doesn't exist",  
    "target": "uuid"  
  }  
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
_links	_links	
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
name	string	Snapshot copy. Valid in POST or PATCH.
owners	array[string]	
snaplock_expiry_time	string	SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time.
state	string	State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.
svm	svm	

Name	Type	Description
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	volume	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a volume Snapshot copy

POST /storage/volumes/{volume.uuid}/snapshots

Introduced In: 9.6

Creates a volume Snapshot copy.

Required properties

- `name` - Name of the Snapshot copy to be created.

Recommended optional properties

- `comment` - Comment associated with the Snapshot copy.
- `expiry_time` - Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.

Related ONTAP commands

- `snapshot create`

Learn more

- [DOC /storage/volumes/{volume.uuid}/snapshots](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
name	string	Snapshot copy. Valid in POST or PATCH.
owners	array[string]	
snaplock_expiry_time	string	SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time.

Name	Type	Description
state	string	State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.
svm	svm	
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	volume	

Example request

```
{
  "comment": "string",
  "create_time": "2019-02-04T19:00:00Z",
  "expiry_time": "2019-02-04T19:00:00Z",
  "name": "this_snapshot",
  "owners": [
    "string"
  ],
  "snaplock_expiry_time": "2019-02-04T19:00:00Z",
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.

Name	Type	Description
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
name	string	Snapshot copy. Valid in POST or PATCH.
owners	array[string]	
snaplock_expiry_time	string	SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time.
state	string	State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.
svm	svm	
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	volume	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a volume Snapshot copy

```
DELETE /storage/volumes/{volume.uuid}/snapshots/{uuid}
```

Introduced In: 9.6

Deletes a Volume Snapshot copy.

Related ONTAP commands

- `snapshot delete`

Learn more

- [DOC /storage/volumes/{volume.uuid}/snapshots](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID

Name	Type	In	Required	Description
uuid	string	path	True	Snapshot copy UUID
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

job_link

Name	Type	Description
_links	_links	
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve volume Snapshot copy details

GET /storage/volumes/{volume.uuid}/snapshots/{uuid}

Introduced In: 9.6

Retrieves details of a specific volume Snapshot copy.

Related ONTAP commands

- `snapshot show`

Learn more

- [DOC /storage/volumes/{volume.uuid}/snapshots](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
uuid	string	path	True	Snapshot copy UUID
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.

Name	Type	Description
name	string	Snapshot copy. Valid in POST or PATCH.
owners	array[string]	
snaplock_expiry_time	string	SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time.
state	string	State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.
svm	svm	
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	volume	

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "comment": "string",
  "create_time": "2019-02-04T19:00:00Z",
  "expiry_time": "2019-02-04T19:00:00Z",
  "name": "this_snapshot",
  "owners": [
    "string"
  ],
  "snaplock_expiry_time": "2019-02-04T19:00:00Z",
  "state": "string",
  "svm": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

svm

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
_links	_links	
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update a volume Snapshot copy

PATCH /storage/volumes/{volume.uuid}/snapshots/{uuid}

Introduced In: 9.6

Updates a Volume Snapshot copy.

Related ONTAP commands

- `snapshot modify`
- `snapshot rename`

Learn more

- [DOC /storage/volumes/{volume.uuid}/snapshots](#)

Parameters

Name	Type	In	Required	Description
volume.uuid	string	path	True	Volume UUID
uuid	string	path	True	Snapshot copy UUID

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.

Name	Type	Description
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
name	string	Snapshot copy. Valid in POST or PATCH.
owners	array[string]	
snaplock_expiry_time	string	SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time.
state	string	State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.
svm	svm	
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	volume	

Example request

```
{
  "comment": "string",
  "create_time": "2019-02-04T19:00:00Z",
  "expiry_time": "2019-02-04T19:00:00Z",
  "name": "this_snapshot",
  "owners": [
    "string"
  ],
  "snaplock_expiry_time": "2019-02-04T19:00:00Z",
  "state": "string",
  "svm": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412",
  "volume": {
    "name": "volume1",
    "uuid": "028baa66-41bd-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default, Error

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

svm

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

volume

Name	Type	Description
name	string	The name of the volume.
uuid	string	Unique identifier for the volume. This corresponds to the instance-uuid that is exposed in the CLI and ONTAPI. It does not change due to a volume move. <ul style="list-style-type: none">• example: 028baa66-41bd-11e9-81d5-00a0986138f7• Introduced in: 9.6

snapshot

The Snapshot copy object represents a point in time Snapshot copy of a volume.

Name	Type	Description
comment	string	A comment associated with the Snapshot copy. This is an optional attribute for POST or PATCH.
create_time	string	Creation time of the Snapshot copy. It is the volume access time when the Snapshot copy was created.

Name	Type	Description
expiry_time	string	The expiry time for the Snapshot copy. This is an optional attribute for POST or PATCH. Snapshot copies with an expiry time set are not allowed to be deleted until the retention time is reached.
name	string	Snapshot copy. Valid in POST or PATCH.
owners	array[string]	
snaplock_expiry_time	string	SnapLock expiry time for the Snapshot copy, if the Snapshot copy is taken on a SnapLock volume. A Snapshot copy is not allowed to be deleted or renamed until the SnapLock ComplianceClock time goes beyond this retention time.
state	string	State of the Snapshot copy. There are cases where some Snapshot copies are not complete. In the "partial" state, the Snapshot copy is consistent but exists only on the subset of the constituents that existed prior to the FlexGroup's expansion. Partial Snapshot copies cannot be used for a Snapshot copy restore operation. A Snapshot copy is in an "invalid" state when it is present in some FlexGroup constituents but not in others. At all other times, a Snapshot copy is valid.
svm	svm	
uuid	string	The UUID of the Snapshot copy in the volume that uniquely identifies the Snapshot copy in that volume.
volume	volume	

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Support

Support overview

Overview

You can use ONTAP support APIs to manage configuration backups, autosupport settings, and event handling.

Configuration backups

Configuration backups are copies of node and cluster settings saved to an external server. Single-node clusters must have configuration backups as protection against corruption of the configuration database. Multi-node clusters back up the configuration automatically between the nodes in the cluster.

Manage AutoSupport configuration

Support AutoSupport endpoint overview

Overview

AutoSupport is the NetApp *call home* mechanism. AutoSupport sends configuration details, status details, and error reporting details to NetApp.

This endpoint supports both GET and PATCH calls. GET is used to retrieve AutoSupport configuration details for the cluster and PATCH is used to modify the AutoSupport configuration of the cluster. You can also use GET calls to check AutoSupport connectivity.

Examples

Configuring 'to' addresses

The following example configures AutoSupport to send emails to 'to' addresses.

```
# The API:
PATCH /support/autosupport

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/autosupport" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"to\": [
\"abc@netapp.com\", \"xyz@netapp.com\" ]}"

# The response:
200 OK
{ }
```

Configuring 'SMTP' transport

The following example configures AutoSupport to use 'SMTP' transport. The default transport is 'HTTPS'.

```
# The API:
PATCH /support/autosupport

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/autosupport" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"transport\": \"smtp\"}"

# The response:
200 OK
{ }
```

Retrieving the AutoSupport configuration

The following example retrieves AutoSupport configuration for the cluster.

```
# The API:
GET /support/autosupport

# The call:
curl -X GET "https://<mgmt-ip>/api/support/autosupport" -H "accept:
application/hal+json"

# The response:
200 OK
{
  "enabled": true,
  "mail_hosts": [
    "mailhost"
  ],
  "from": "Postmaster",
  "to": [
    "abc@netapp.com",
    "xyz@netapp.com"
  ],
  "contact_support": true,
  "transport": "smtp",
  "proxy_url": "",
  "is_minimal": false,
  "_links": {
    "self": {
      "href": "/api/support/autosupport"
    }
  }
}
```

Retrieving AutoSupport connectivity issues

The following example retrieves AutoSupport connectivity issues for the cluster. The `fields=issues` parameter must be specified, for the response to return connectivity issues. The `corrective_action` section might contain commands which needs to be executed on the ONTAP CLI.

Note that the connectivity check can take up to 10 seconds to complete.

```
# The API:
GET /support/autosupport
```

```
# The call:
curl -X GET "https://<mgmt-ip>/api/support/autosupport?fields=issues" -H
"accept: application/hal+json"

# The response:
200 OK
{
  "issues": [
    {
      "node": {
        "name": "node3",
        "uuid": "0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/0ecfd0a6-f1b3-11e8-9d9f-
005056bbaadc"
          }
        }
      },
      "issue": {
        "message": "SMTP connectivity check failed for destination:
mailhost. Error: Could not resolve host - 'mailhost'",
        "code": "53149746"
      },
      "corrective_action": {
        "message": "Check the hostname of the SMTP server",
        "code": "53149746"
      }
    },
    {
      "node": {
        "name": "node3",
        "uuid": "0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/0ecfd0a6-f1b3-11e8-9d9f-
005056bbaadc"
          }
        }
      },
      "issue": {
        "message": "AutoSupport OnDemand is disabled when \"-transport\" is
not set to \"https\".",
        "code": "53149740"
      },
      "corrective_action": {

```

```

    "message": "Run \"system node autosupport modify -transport https
-node <node name>\" to set \"-transport\" to \"https\".",
    "code": "53149740"
  }
}
],
"_links": {
  "self": {
    "href": "/api/support/autosupport"
  }
}
}
}

```

Retrieving AutoSupport configuration and connectivity issues

The following example retrieves AutoSupport configuration and connectivity issues on the cluster. Use `fields=*,issues` parameter to return both configuration and connectivity issues.

```

# The API:
GET /support/autosupport

# The call:
curl -X GET "https://<mgmt-ip>/api/support/autosupport?fields=%2Cissues"
-H "accept: application/hal+json"

# The response:
200 OK
{
  "enabled": true,
  "mail_hosts": [
    "mailhost"
  ],
  "from": "Postmaster",
  "to": [
    "abc@netapp.com",
    "xyz@netapp.com"
  ],
  "contact_support": true,
  "transport": "smtp",
  "proxy_url": "",
  "is_minimal": false,
  "issues": [
    {
      "node": {

```

```

    "name": "node3",
    "uuid": "0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc"
      }
    },
    "issue": {
      "message": "SMTP connectivity check failed for destination: mailhost. Error: Could not resolve host - 'mailhost'",
      "code": "53149746"
    },
    "corrective_action": {
      "message": "Check the hostname of the SMTP server",
      "code": "53149746"
    }
  },
  {
    "node": {
      "name": "node3",
      "uuid": "0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/0ecfd0a6-f1b3-11e8-9d9f-005056bbaadc"
        }
      }
    },
    "issue": {
      "message": "AutoSupport OnDemand is disabled when \"-transport\" is not set to \"https\".",
      "code": "53149740"
    },
    "corrective_action": {
      "message": "Run \"system node autosupport modify -transport https -node <node name>\" to set \"-transport\" to \"https\".",
      "code": "53149740"
    }
  }
],
"_links": {
  "self": {
    "href": "/api/support/autosupport"
  }
}

```

```
}  
}
```

Retrieve the AutoSupport configuration

GET /support/autosupport

Introduced In: 9.6

Retrieves the AutoSupport configuration of the cluster and if requested, returns connectivity issues with the AutoSupport configuration.

Important note:

- The **issues** field consists of a list of objects containing details of the node that has a connectivity issue, the issue description, and corrective action you can take to address the issue. When not empty, this indicates a connection issue to the **HTTP/S**, **SMTP**, or **AutoSupport On Demand** server.

Expensive properties

There is an added cost to retrieving values for these properties. They are not included by default in GET results and must be explicitly requested using the `fields` query parameter. See [Requesting specific fields](#) to learn more.

- `issues`

Related ONTAP commands

- `system node autosupport show -instance`
- `system node autosupport check show-details`

Learn more

- [DOC /support/autosupport](#)

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
contact_support	boolean	Specifies whether to send the AutoSupport messages to vendor support.
enabled	boolean	Specifies whether the AutoSupport daemon is enabled. When this setting is disabled, delivery of all AutoSupport messages is turned off.

Name	Type	Description
from	string	<p>The e-mail address from which the AutoSupport messages are sent. To generate node-specific 'from' addresses, enable '-node-specific-from' parameter via ONTAP CLI.</p> <ul style="list-style-type: none"> example: postmaster@example.com format: email Introduced in: 9.6
is_minimal	boolean	Specifies whether the system information is collected in compliant form, to remove private data or in complete form, to enhance diagnostics.
issues	array[autosupport_issues]	A list of nodes in the cluster with connectivity issues to HTTP/SMTP/AOD AutoSupport destinations along with the corresponding error descriptions and corrective actions.
mail_hosts	array[string]	The names of the mail servers used to deliver AutoSupport messages via SMTP.
partner_addresses	array[string]	The list of partner addresses.
proxy_url	string	Proxy server for AutoSupport message delivery via HTTP/S. Optionally specify a username/password for authentication with the proxy server.
to	array[string]	The e-mail addresses to which the AutoSupport messages are sent.
transport	string	The name of the transport protocol used to deliver AutoSupport messages.

Example response

```
{
  "contact_support": 1,
  "enabled": 1,
  "from": "<a href="
mailto:postmaster@example.com">postmaster@example.com</a>",
  "is_minimal": 1,
  "issues": [
    {
      "corrective_action": {
        "code": "53149746",
        "message": "Check the hostname of the SMTP server"
      },
      "issue": {
        "code": "53149746",
        "message": "SMTP connectivity check failed for destination:
mailhost. Error: Could not resolve host - 'mailhost'"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "mail_hosts": [
    "mailhost1.example.com",
    "mailhost2.example.com"
  ],
  "partner_addresses": [
    "<a href="mailto:user1@partner.com">user1@partner.com</a>",
    "<a href="mailto:user2@partner.com">user2@partner.com</a>"
  ],
  "proxy_url": "https://proxy.company.com",
  "to": [
    "<a href="mailto:user1@example.com">user1@example.com</a>",
    "<a href="mailto:user2@example.com">user2@example.com</a>"
  ],
  "transport": "smtp"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

autosupport_connectivity_corrective_action

Name	Type	Description
code	string	Corrective action code
message	string	<p>Corrective action message. The corrective action might contain commands which needs to be executed on the ONTAP CLI.</p> <ul style="list-style-type: none">• example: Check the hostname of the SMTP server• readOnly: 1• Introduced in: 9.6

autosupport_connectivity_issue

Name	Type	Description
code	string	Error code
message	string	Error message

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

autosupport_issues

Name	Type	Description
corrective_action	autosupport_connectivity_corrective_action	
issue	autosupport_connectivity_issue	
node	node	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the AutoSupport configuration for an entire cluster

PATCH /support/autosupport

Introduced In: 9.6

Updates the AutoSupport configuration for the entire cluster.

Related ONTAP commands

- `system node autosupport modify`

Learn more

- [DOC /support/autosupport](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
contact_support	boolean	Specifies whether to send the AutoSupport messages to vendor support.
enabled	boolean	Specifies whether the AutoSupport daemon is enabled. When this setting is disabled, delivery of all AutoSupport messages is turned off.
from	string	<p>The e-mail address from which the AutoSupport messages are sent. To generate node-specific 'from' addresses, enable '-node-specific-from' parameter via ONTAP CLI.</p> <ul style="list-style-type: none"> • example: postmaster@example.com • format: email • Introduced in: 9.6
is_minimal	boolean	Specifies whether the system information is collected in compliant form, to remove private data or in complete form, to enhance diagnostics.
issues	array[autosupport_issues]	A list of nodes in the cluster with connectivity issues to HTTP/SMTP/AOD AutoSupport destinations along with the corresponding error descriptions and corrective actions.
mail_hosts	array[string]	The names of the mail servers used to deliver AutoSupport messages via SMTP.

Name	Type	Description
partner_addresses	array[string]	The list of partner addresses.
proxy_url	string	Proxy server for AutoSupport message delivery via HTTP/S. Optionally specify a username/password for authentication with the proxy server.
to	array[string]	The e-mail addresses to which the AutoSupport messages are sent.
transport	string	The name of the transport protocol used to deliver AutoSupport messages.

Example request

```
{
  "contact_support": 1,
  "enabled": 1,
  "from": "<a href="
mailto:postmaster@example.com">postmaster@example.com</a>",
  "is_minimal": 1,
  "issues": [
    {
      "corrective_action": {
        "code": "53149746",
        "message": "Check the hostname of the SMTP server"
      },
      "issue": {
        "code": "53149746",
        "message": "SMTP connectivity check failed for destination:
mailhost. Error: Could not resolve host - 'mailhost'"
      },
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ],
  "mail_hosts": [
    "mailhost1.example.com",
    "mailhost2.example.com"
  ],
  "partner_addresses": [
    "<a href="mailto:user1@partner.com">user1@partner.com</a>",
    "<a href="mailto:user2@partner.com">user2@partner.com</a>"
  ],
  "proxy_url": "https://proxy.company.com",
  "to": [
    "<a href="mailto:user1@example.com">user1@example.com</a>",
    "<a href="mailto:user2@example.com">user2@example.com</a>"
  ],
  "transport": "smtp"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8650862	The SMTP mail host provided cannot be empty
8650863	A maximum of 5 SMTP mail hosts can be provided
8650864	A maximum of 5 email addresses can be provided
8650865	A maximum of 5 partner email addresses can be provided
53149727	The proxy URI provided is invalid
53149728	The mailhost URI provided is invalid

Definitions

See Definitions

autosupport_connectivity_corrective_action

Name	Type	Description
code	string	Corrective action code
message	string	<p>Corrective action message. The corrective action might contain commands which needs to be executed on the ONTAP CLI.</p> <ul style="list-style-type: none">• example: Check the hostname of the SMTP server• readOnly: 1• Introduced in: 9.6

autosupport_connectivity_issue

Name	Type	Description
code	string	Error code
message	string	Error message

href

Name	Type	Description
href	string	

_links

node

Name	Type	Description
name	string	
uuid	string	

autosupport_issues

Name	Type	Description
corrective_action	autosupport_connectivity_corrective_action	
issue	autosupport_connectivity_issue	
node	node	

autosupport

Name	Type	Description
contact_support	boolean	Specifies whether to send the AutoSupport messages to vendor support.
enabled	boolean	Specifies whether the AutoSupport daemon is enabled. When this setting is disabled, delivery of all AutoSupport messages is turned off.
from	string	<p>The e-mail address from which the AutoSupport messages are sent. To generate node-specific 'from' addresses, enable '-node-specific-from' parameter via ONTAP CLI.</p> <ul style="list-style-type: none"> • example: postmaster@example.com • format: email • Introduced in: 9.6
is_minimal	boolean	Specifies whether the system information is collected in compliant form, to remove private data or in complete form, to enhance diagnostics.
issues	array[autosupport_issues]	A list of nodes in the cluster with connectivity issues to HTTP/SMTP/AOD AutoSupport destinations along with the corresponding error descriptions and corrective actions.
mail_hosts	array[string]	The names of the mail servers used to deliver AutoSupport messages via SMTP.
partner_addresses	array[string]	The list of partner addresses.

Name	Type	Description
proxy_url	string	Proxy server for AutoSupport message delivery via HTTP/S. Optionally specify a username/password for authentication with the proxy server.
to	array[string]	The e-mail addresses to which the AutoSupport messages are sent.
transport	string	The name of the transport protocol used to deliver AutoSupport messages.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage AutoSupport messages

Support AutoSupport messages endpoint overview

Overview

Use this API to invoke and retrieve AutoSupport messages from the nodes in the cluster.

This API supports POST and GET calls. Use a POST call to invoke AutoSupport and a GET call to retrieve AutoSupport messages.

Examples

Invoking an AutoSupport on all nodes in the cluster

The following example invokes an AutoSupport on every node in the cluster.

Note that AutoSupport is invoked on all nodes in the cluster if the `node` parameter is omitted. Also, note that the `subject` line is the same when invoking on all nodes.

By default, the response is an empty object. If `return_records=true` is passed in the request, the response includes information about the node and the index of the invoked AutoSupport message.

```
# The API:
POST /support/autosupport/messages

# The call:
curl -X POST "https://<mgmt-
ip>/api/support/autosupport/messages?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{ \"
message\": \"test_msg\", \"type\": \"all\"}"

# The response:
201 CREATED
{
  "num_records": 2,
  "records": [
    {
      "index": 4,
      "node": {
        "name": "node1",
        "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/4"
        }
      }
    },
    {
      "index": 2,
      "node": {
```

```

    "name": "node2",
    "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-005056bb5cab"
      }
    },
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-b186-005056bb5cab/2"
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/support/autosupport/messages/"
    }
  }
}

```

Invoking an AutoSupport on a single node

The following examples invoke an AutoSupport on a single node in the cluster.

Note that AutoSupport is invoked on all nodes in the cluster if the `node` parameter is omitted. You can specify the node-name with either `node` or `node.name` parameter. You can also specify UUID of the node with the `node.uuid` parameter.

By default, the response is an empty object. If `return_records=true` is passed in the request, the response includes information about the node and the index of the invoked AutoSupport message.

```

# The API:
POST /support/autosupport/messages

# The call:
curl -X POST "https://<mgmt-
ip>/api/support/autosupport/messages?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
  \"message\": \"test_msg\", \"type\": \"test\", \"node\": { \"name\":
  \"node1\" } }"

# The response:
201 CREATED
{
  "num_records": 1,
  "records": [
    {
      "index": 8,
      "node": {
        "name": "node1",
        "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/8"
        }
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/support/autosupport/messages/"
    }
  }
}

```

```
# The call:
curl -X POST "https://<mgmt-
ip>/api/support/autosupport/messages?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
\"message\": \"test_msg\", \"type\": \"test\", \"node.name\": \"node2\"}"

# The response:
201 CREATED
{
  "num_records": 1,
  "records": [
    {
      "index": 4,
      "node": {
        "name": "node2",
        "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-
005056bb5cab"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-
b186-005056bb5cab/4"
        }
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/support/autosupport/messages/"
    }
  }
}
```

```
# The call:
curl -X POST "https://<mgmt-
ip>/api/support/autosupport/messages?return_records=true" -H "accept:
application/hal+json" -H "Content-Type: application/json" -d "{
  \"message\": \"test_msg\", \"type\": \"test\", \"node.uuid\": \"092e0298-
f250-11e8-9a05-005056bb6666\"}"

# The response:
201 CREATED
{
  "num_records": 1,
  "records": [
    {
      "index": 5,
      "node": {
        "name": "node1",
        "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/5"
        }
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/support/autosupport/messages/"
    }
  }
}
```

Retrieving AutoSupport messages from all nodes in the cluster

The following example retrieves AutoSupport messages from every node in the cluster.

Note that if the *fields=** parameter is not specified, only node, index, and destination fields are returned.

Filters can be added on the fields to limit the results.


```
# The API:
GET /support/autosupport/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/support/autosupport/messages?fields=*&return_timeout=15" -H
"accept: application/hal+json"

# The response:
200 OK
{
  "records": [
    {
      "node": {
        "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
          }
        }
      },
      "index": 1,
      "destination": "smtp",
      "subject": "USER_TRIGGERED (TEST:test_msg)",
      "state": "ignore",
      "generated_on": "2019-03-28T10:18:04-04:00",
      "_links": {
        "self": {
          "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/1/smtp"
        }
      }
    },
    {
      "node": {
        "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
          }
        }
      }
    }
  ]
}
```

```

    },
    "index": 1,
    "destination": "http",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "sent_successful",
    "generated_on": "2019-03-28T10:18:04-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-9a05-005056bb6666/1/http"
      }
    }
  },
  {
    "node": {
      "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
      "name": "node1",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-005056bb6666"
        }
      }
    },
    "index": 1,
    "destination": "noteto",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "ignore",
    "generated_on": "2019-03-28T10:18:04-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/092e0298-f250-11e8-9a05-005056bb6666/1/noteto"
      }
    }
  },
  {
    "node": {
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-005056bb5cab"
        }
      }
    }
  }
}

```

```

    },
    "index": 1,
    "destination": "smtp",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "ignore",
    "generated_on": "2019-03-28T10:18:06-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-b186-005056bb5cab/1/smtp"
      }
    }
  },
  {
    "node": {
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-005056bb5cab"
        }
      }
    },
    "index": 1,
    "destination": "http",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "sent_successful",
    "generated_on": "2019-03-28T10:18:06-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-b186-005056bb5cab/1/http"
      }
    }
  },
  {
    "node": {
      "uuid": "e47d2630-f250-11e8-b186-005056bb5cab",
      "name": "node2",
      "_links": {
        "self": {
          "href": "/api/cluster/nodes/e47d2630-f250-11e8-b186-005056bb5cab"
        }
      }
    }
  }
}

```

```

    },
    "index": 1,
    "destination": "noteto",
    "subject": "USER_TRIGGERED (TEST:test_msg)",
    "state": "ignore",
    "generated_on": "2019-03-28T10:18:06-04:00",
    "_links": {
      "self": {
        "href": "/api/support/autosupport/messages/e47d2630-f250-11e8-b186-005056bb5cab/1/noteto"
      }
    }
  },
],
"num_records": 6,
"_links": {
  "self": {
    "href": "/api/support/autosupport/messages?fields=*&return_timeout=15"
  }
}
}

```

Retrieving AutoSupport messages from a specific node and has 'sent_succesful' state

The following example retrieves AutoSupport messages from a specific node in the cluster.

Note that if the `fields=*` parameter is not specified, only `node`, `index`, and `destination` fields are returned.

This example uses a filter on the `node.name` and `state` fields. You can add filters to any fields in the response.

```

# The API:
GET /support/autosupport/messages

# The call:
curl -X GET "https://<mgmt-
ip>/api/support/autosupport/messages?node.name=node1&state=sent_successful
&fields=*&return_timeout=15" -H "accept: application/hal+json"

# The response:
200 OK
{
  "records": [
    {
      "node": {
        "uuid": "092e0298-f250-11e8-9a05-005056bb6666",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/092e0298-f250-11e8-9a05-
005056bb6666"
          }
        }
      },
      "index": 1,
      "destination": "http",
      "subject": "USER_TRIGGERED (TEST:test_msg)",
      "state": "sent_successful",
      "generated_on": "2019-03-28T10:18:04-04:00",
      "_links": {
        "self": {
          "href": "/api/support/autosupport/messages/092e0298-f250-11e8-
9a05-005056bb6666/1/http"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href":
"/api/support/autosupport/messages?node.name=node1&state=sent_successful&f
ields=*&return_timeout=15"
    }
  }
}

```

Retrieve AutoSupport message history

GET /support/autosupport/messages

Introduced In: 9.6

Retrieves AutoSupport message history from all nodes in the cluster.

There can be a short delay on invoked AutoSupport messages showing in history, dependent on processing of other AutoSupports in the queue.

Related ONTAP commands

- `system node autosupport history show`

Learn more

- [DOC /support/autosupport/messages](#)

Parameters

Name	Type	In	Required	Description
node.uuid	string	query	False	Filter by node.uuid
node.name	string	query	False	Filter by node.name
destination	string	query	False	Filter by destination
state	string	query	False	Filter by state
generated_on	string	query	False	Filter by generated_on
index	integer	query	False	Filter by index
subject	string	query	False	Filter by subject
error.message	string	query	False	Filter by error.message
error.code	integer	query	False	Filter by error.code
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False
Specify the fields to return.	max_records	integer	query	False
Limit the number of records returned.	return_timeout	integer	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[autosupport_message]	List of messages invoked on the node

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "3",
  "records": [
    {
      "destination": "http",
      "error": {
        "code": "53149746",
        "message": "Could not resolve host: test.com"
      },
      "generated_on": "2019-03-25T17:30:04-04:00",
      "index": "9",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "state": "sent_successful",
      "subject": "WEEKLY_LOG"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

error

Last error during delivery attempt. Empty if "status=sent-successful".

Name	Type	Description
code	integer	Error code
message	string	Error message

_links

Name	Type	Description
self	href	

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

autosupport_message

Name	Type	Description
destination	string	Destination for the AutoSupport
error	error	Last error during delivery attempt. Empty if "status=sent-successful".

Name	Type	Description
generated_on	string	Date and Time of AutoSupport generation in ISO-8601 format
index	integer	Sequence number of the AutoSupport
node	node	
state	string	State of AutoSupport delivery
subject	string	Subject line for the AutoSupport

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create and send an AutoSupport message

POST /support/autosupport/messages

Introduced In: 9.6

Creates and sends an AutoSupport message with the provided input parameters.

Important note:

- By default, the response is an empty object. If `return_records=true` is passed in the request, the response includes information about the node and the index of the invoked AutoSupport message.

Recommended optional properties

- `message` - Message included in the AutoSupport subject. Use this to identify the generated AutoSupport message.

Default property values

If not specified in POST, the following are the default property values:

- `type` - *all*
- `node.name` or `node.uuid` - Not specifying these properties invokes AutoSupport on all nodes in the cluster.

Related ONTAP commands

- `system node autosupport invoke`

Learn more

- [DOC /support/autosupport/messages](#)

Parameters

Name	Type	In	Required	Description
<code>return_records</code>	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
<code>destination</code>	string	Destination for the AutoSupport
<code>error</code>	error	Last error during delivery attempt. Empty if "status=sent-successful".
<code>generated_on</code>	string	Date and Time of AutoSupport generation in ISO-8601 format
<code>index</code>	integer	Sequence number of the AutoSupport
<code>message</code>	string	Message included in the AutoSupport subject
<code>node</code>	node	

Name	Type	Description
state	string	State of AutoSupport delivery
subject	string	Subject line for the AutoSupport
type	string	Type of AutoSupport collection to issue
uri	string	Alternate destination for the AutoSupport

Example request

```
{
  "destination": "http",
  "error": {
    "code": "53149746",
    "message": "Could not resolve host: test.com"
  },
  "generated_on": "2019-03-25T17:30:04-04:00",
  "index": "9",
  "message": "invoked_test_autosupport_rest",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "state": "sent_successful",
  "subject": "WEEKLY_LOG",
  "type": "test",
  "uri": "http://1.2.3.4/delivery_uri"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "num_records": "3",
  "records": [
    {
      "index": "9",
      "node": {
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      }
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
53149748	The destination URI provided for the invoked AutoSupport is invalid
655294464	The message parameter is not supported with performance AutoSupports

Definitions

See Definitions

error

Last error during delivery attempt. Empty if "status=sent-successful".

Name	Type	Description
code	integer	Error code
message	string	Error message

href

Name	Type	Description
href	string	

_links

node

Name	Type	Description
name	string	
uuid	string	

autosupport_message

Name	Type	Description
destination	string	Destination for the AutoSupport
error	error	Last error during delivery attempt. Empty if "status=sent-successful".
generated_on	string	Date and Time of AutoSupport generation in ISO-8601 format
index	integer	Sequence number of the AutoSupport
message	string	Message included in the AutoSupport subject
node	node	
state	string	State of AutoSupport delivery
subject	string	Subject line for the AutoSupport

Name	Type	Description
type	string	Type of AutoSupport collection to issue
uri	string	Alternate destination for the AutoSupport

records

Name	Type	Description
index	integer	Sequence number of the generated AutoSupport
node	node	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster configuration backup information

Support configuration-backup endpoint overview

Overview

This API retrieves the current settings for the configuration and updates configuration backup settings. The GET operation retrieves the current settings for the configuration and the PATCH operation updates the configuration backup settings.

Examples

These examples show how to retrieve and update the configuration backup settings.

Retrieving the configuration backup settings

```
# The API:
/api/support/configuration-backup

# The call:
curl -X GET "https://<mgmt-ip>/api/support/configuration-backup" -H
"accept: application/hal+json"

# The response:
{
  "url": "http://10.224.65.198/backups",
  "username": "me",
  "_links": {
    "self": {
      "href": "/api/support/configuration-backup"
    }
  }
}
```

Updating the configuration backup settings

```
# The API:
/api/support/configuration-backup

# The call:
curl -X PATCH "https://<mgmt-ip>/api/support/configuration-backup" -H
"accept: application/hal+json"

# The body:
{
  "url": "https://10.224.65.198/new_backups",
  "username": "new_me",
  "password": "new_pass",
  "validate_certificate": "true"
}

# The response header:
HTTP/1.1 200 OK
Date: Tue, 05 Jun 2018 18:17:48 GMT
Server: libzapid-httpd
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/hal+json
```

Retrieve cluster configuration backup information

GET /support/configuration-backup

Introduced In: 9.6

Retrieves the cluster configuration backup information.

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
password	string	
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	

Name	Type	Description
validate_certificate	boolean	Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default.

Example response

```
{
  "password": "yourpassword",
  "url": "http://10.224.65.198/backups",
  "username": "me"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update cluster configuration backup information

PATCH /support/configuration-backup

Introduced In: 9.6

Updates the cluster configuration backup information.

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
password	string	

Name	Type	Description
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	
validate_certificate	boolean	Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default.

Example request

```
{
  "password": "yourpassword",
  "url": "http://10.224.65.198/backups",
  "username": "me"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
39387137	Invalid URL
39387138	Invalid URL Scheme
5963792	Using HTTPS for URL requires an effective cluster version of 9.7.0 or later.

Error Code	Description
5963794	Validating the certificate is supported for the HTTPS protocol only.

Definitions

See Definitions

configuration_backup

Name	Type	Description
password	string	
url	string	An external backup location for the cluster configuration. This is mostly required for single node clusters where node and cluster configuration backups cannot be copied to other nodes in the cluster.
username	string	
validate_certificate	boolean	Use this parameter with the value "true" to validate the digital certificate of the remote server. Digital certificate validation is available only when the HTTPS protocol is used in the URL; it is disabled by default.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage configuration backups

Support configuration-backup backups endpoint overview

Overview

This API supports creating, deleting and retrieving configuration backup files.

Configuration backups can be 'cluster' or 'node' type. A 'cluster' backup contains cluster-wide configuration in addition to the configuration of each node in the cluster. A 'node' backup contains only node-specific configuration such as configuration files on the root volume and the boot variables. For creating a cluster backup, a cluster-wide job is queued. For creating a node backup, a private job local to the node is queued. In addition to the backups created using this API, ONTAP creates configuration backups automatically based on job schedules. This API supports creating configuration backups on demand only. It supports deleting and retrieving configuration backups that are created automatically or on demand.

For information on configuration backup settings for automatically created backups, see [DOC /support/configuration-backup](#)

Examples

Retrieving a list of configuration backup files

```
# The API:
/api/support/configuration-backup/backups

# The call:
curl -X GET "https://<mgmt-ip>/api/support/configuration-backup/backups"
-H "accept: application/hal+json"

# The response:
{
  "records": [
    {
      "node": {
        "uuid": "5cafe0f6-499f-11e9-b644-005056bbcf93",
        "name": "node1",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/5cafe0f6-499f-11e9-b644-005056bbcf93"
          }
        }
      },
      "name": "backup1.7z",
      "_links": {
        "self": {
          "href": "/api/support/configuration-backup/backups/5cafe0f6-499f-11e9-b644-005056bbcf93/backup1.7z"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/support/configuration-backup/backups"
    }
  }
}
```

Retrieving details of the specified configuration backup file

```
# The API:
/api/support/configuration-backup/backups/{node.uuid}/{name}

# The call:
curl -X GET "https://<mgmt-ip>/api/support/configuration-
backup/backups/bc2f15d0-8b93-11e9-90e9-005056bb6a30/backup1.7z" -H
"accept: application/hal+json"

# The response:
{
  "node": {
    "uuid": "bc2f15d0-8b93-11e9-90e9-005056bb6a30",
    "name": "node1",
    "_links": {
      "self": {
        "href": "/api/cluster/nodes/bc2f15d0-8b93-11e9-90e9-005056bb6a30"
      }
    }
  },
  "name": "backup1.7z",
  "type": "cluster",
  "time": "2019-06-10T13:35:06-04:00",
  "size": 6058408,
  "backup_nodes": [
    {
      "name": "node1"
    },
    {
      "name": "node2"
    }
  ],
  "version": "9.7.0",
  "auto": false,
  "download_link": "https://10.224.66.113/backups/backup1.7z",
  "_links": {
    "self": {
      "href": "/api/support/configuration-backup/backups/bc2f15d0-8b93-11e9-
90e9-005056bb6a30/backup1.7z"
    }
  }
}
```

Creating a configuration backup file

```
# The API:
/api/support/configuration-backup/backups

# The call:
curl -X POST "https://<mgmt-ip>/api/support/configuration-backup/backups"
-H "accept: application/hal+json" -d "@create_backup.txt"
create_backup.txt:
{
  "node":
    {
      "uuid": "ac13c636-4fc9-11e9-94c2-005056bb2516",
      "name": "node1"
    },
  "name": "backup3.7z"
}

# The response header:
HTTP/1.1 202 Accepted
Date: Tue, 26 Mar 2019 14:26:24 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Location: /api/support/configuration-backup/backups/ac13c636-4fc9-11e9-94c2-005056bb2516/backup3.7z
Content-Length: 189
Content-Type: application/hal+json
{
  "job": {
    "uuid": "22acfb68-4fd3-11e9-94c2-005056bb2516",
    "_links": {
      "self": {
        "href": "/api/cluster/jobs/22acfb68-4fd3-11e9-94c2-005056bb2516"
      }
    }
  }
}
```

Deleting a configuration backup file

```
# The API:
/api/support/configuration-backup/backups/{node.uuid}/{name}

# The call:
curl -X DELETE "https://<mgmt-ip>/api/support/configuration-backup/backups/5cafe0f6-499f-11e9-b644-005056bbcf93/backup1.7z" -H "content-type: application/json"

# The response header:
HTTP/1.1 200 OK
Date: Tue, 26 Mar 2019 14:32:23 GMT
Server: libzapid-httpd
X-Content-Type-Options: nosniff
Cache-Control: no-cache,no-store,must-revalidate
Content-Length: 3
Content-Type: application/hal+json
```

Retrieve configuration backup files

GET /support/configuration-backup/backups

Introduced In: 9.7

Retrieves a list of configuration backup files.

Related ONTAP commands

- `system configuration backup show`

Parameters

Name	Type	In	Required	Description
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
backup_nodes.name	string	query	False	Filter by backup_nodes.name
auto	boolean	query	False	Filter by auto
time	string	query	False	Filter by time

Name	Type	In	Required	Description
download_link	string	query	False	Filter by download_link
version	string	query	False	Filter by version
name	string	query	False	Filter by name
size	integer	query	False	Filter by size
type	string	query	False	Filter by type
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[configuration_backup_file]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "backup_nodes": [
        {
          "name": "string"
        }
      ],
      "download_link": "https://10.224.65.198/backups/backup_file.7z",
      "name": "backup_file.7z",
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "size": "4787563",
      "time": "2019-02-04T13:33:48-05:00",
      "type": "string",
      "version": "9.7.0"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

backup_node

The node name in the backup.

Name	Type	Description
name	string	

_links

Name	Type	Description
self	href	

node

The node that owns the configuration backup.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

configuration_backup_file

The configuration backup file.

Name	Type	Description
auto	boolean	Indicates if the backup was created automatically.
backup_nodes	array[backup_node]	The list of nodes included in the backup.

Name	Type	Description
download_link	string	The link to download the backup file.
name	string	The backup name.
node	node	The node that owns the configuration backup.
size	integer	The size of the backup in bytes.
time	string	The backup creation time.
type	string	The backup type.
version	string	The software version.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a configuration backup

POST /support/configuration-backup/backups

Introduced In: 9.7

Creates a configuration backup. The required backup file name must end with .7z extension.

Required properties

- `node.uuid` or `node.name` - The node UUID or node name on which the configuration backup will be created.
- `name` - The backup file name

Related ONTAP commands

- `system configuration backup create`

Parameters

Name	Type	In	Required	Description
<code>return_timeout</code>	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Name	Type	In	Required	Description
return_records	boolean	query	False	<p>The default is false. If set to true, the records are returned.</p> <ul style="list-style-type: none"> • Default value:

Request Body

Name	Type	Description
auto	boolean	Indicates if the backup was created automatically.
backup_nodes	array[backup_node]	The list of nodes included in the backup.
download_link	string	The link to download the backup file.
name	string	The backup name.
node	node	The node that owns the configuration backup.
size	integer	The size of the backup in bytes.
time	string	The backup creation time.
type	string	The backup type.
version	string	The software version.

Example request

```
{
  "backup_nodes": [
    {
      "name": "string"
    }
  ],
  "download_link": "https://10.224.65.198/backups/backup_file.7z",
  "name": "backup_file.7z",
  "node": {
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": "4787563",
  "time": "2019-02-04T13:33:48-05:00",
  "type": "string",
  "version": "9.7.0"
}
```

Response

Status: 202, Accepted

Name	Type	Description
job	job_link	

Example response

```
{
  "job": {
    "uuid": "string"
  }
}
```

Error

Status: Default

Error Code	Description
5963818	Failed to queue private job.
5963819	Failed to queue cluster job.
5963871	Requested backup name is missing .7z suffix.

Definitions

See Definitions

backup_node

The node name in the backup.

Name	Type	Description
name	string	

href

Name	Type	Description
href	string	

_links

node

The node that owns the configuration backup.

Name	Type	Description
name	string	
uuid	string	

configuration_backup_file

The configuration backup file.

Name	Type	Description
auto	boolean	Indicates if the backup was created automatically.
backup_nodes	array[backup_node]	The list of nodes included in the backup.
download_link	string	The link to download the backup file.
name	string	The backup name.
node	node	The node that owns the configuration backup.
size	integer	The size of the backup in bytes.
time	string	The backup creation time.

Name	Type	Description
type	string	The backup type.
version	string	The software version.

job_link

Name	Type	Description
uuid	string	The UUID of the asynchronous job that is triggered by a POST, PATCH, or DELETE operation.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Delete a configuration backup

DELETE /support/configuration-backup/backups/{node.uuid}/{name}

Introduced In: 9.7

Deletes a configuration backup.

Related ONTAP commands

- `system configuration backup delete`

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	UUID of the node that owns the configuration backup.
name	string	path	True	Name of the configuration backup to be deleted.

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5963826	Failed to delete backup file.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve configuration backup file details

GET /support/configuration-backup/backups/{node.uuid}/{name}

Introduced In: 9.7

Retrieves details of the specified configuration backup file.

Related ONTAP commands

- `system configuration backup show`

Parameters

Name	Type	In	Required	Description
node.uuid	string	path	True	UUID of the node that owns the configuration backup.
name	string	path	True	Name of the configuration backup.
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
auto	boolean	Indicates if the backup was created automatically.
backup_nodes	array[backup_node]	The list of nodes included in the backup.
download_link	string	The link to download the backup file.
name	string	The backup name.
node	node	The node that owns the configuration backup.
size	integer	The size of the backup in bytes.
time	string	The backup creation time.

Name	Type	Description
type	string	The backup type.
version	string	The software version.

Example response

```
{
  "backup_nodes": [
    {
      "name": "string"
    }
  ],
  "download_link": "https://10.224.65.198/backups/backup_file.7z",
  "name": "backup_file.7z",
  "node": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "node1",
    "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
  },
  "size": "4787563",
  "time": "2019-02-04T13:33:48-05:00",
  "type": "string",
  "version": "9.7.0"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
5963777	Configuration backup file does not exist.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

backup_node

The node name in the backup.

Name	Type	Description
name	string	

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

node

The node that owns the configuration backup.

Name	Type	Description
_links	_links	
name	string	
uuid	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Manage EMS configuration

Support EMS endpoint overview

Overview

The Event Management System (EMS) collects and processes events, and sends notification of the events through various reporting mechanisms. The following endpoints defined under '/support/ems', allow you to query a list of observed events, and configure which events you handle and how you are notified:

- /support/ems
- /support/ems/events
- /support/ems/messages
- /support/ems/filters
- /support/ems/filters/{name}/rules
- /support/ems/filters/{name}/rules/{index}
- /support/ems/destinations
- /support/ems/destinations/{name}

Examples

Configuring an email destination

The following example configures EMS to send a support email when a WAFL event is observed with an error severity.

Configure the system-wide email parameters

```
# API
PATCH /support/ems

# JSON Body
{
  "mail_from": "administrator@mycompany.com",
  "mail_server": "smtp@mycompany.com"
}

# Response
200 OK
```


Configuring a filter with an enclosed rule

```
# API
POST /support/ems/filters

# JSON Body
{
  "name": "critical-wafl",
  "rules": [
    {
      "index": 1,
      "type": "include",
      "message_criteria": {
        "name_pattern": "wafl.*",
        "severities": "emergency,error,alert"
      }
    }
  ]
}

# Response
201 Created
```

Setting up an email destination

```
# API
POST /support/ems/destinations

# JSON Body
{
  "name": "Technician_Email",
  "type": "email",
  "destination": "technician@mycompany.com",
  "filters": [
    { "name" : "critical-wafl" }
  ]
}

# Response
201 Created
```

Retrieve an EMS configuration

GET /support/ems

Introduced In: 9.6

Retrieves the EMS configuration.

Related ONTAP commands

- `event config show`

Parameters

Name	Type	In	Required	Description
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc
desc] direction. Default direction is 'asc' for ascending.	fields	array[string]	query	False

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
mail_from	string	Mail from
mail_server	string	Mail server (SMTP)
proxy_password	string	Password for HTTP/HTTPS proxy
proxy_url	string	HTTP/HTTPS proxy URL
proxy_user	string	User name for HTTP/HTTPS proxy
pubsub_enabled	boolean	Is Publish/Subscribe Messaging Enabled?

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "mail_from": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "mail_server": "<a href="
mailto:mail@mycompany.com">mail@mycompany.com</a>",
  "proxy_password": "password",
  "proxy_url": "https://proxyserver.mycompany.com",
  "proxy_user": "proxy_user",
  "pubsub_enabled": 1
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update an EMS configuration

PATCH /support/ems

Introduced In: 9.6

Updates the EMS configuration.

Related ONTAP commands

- `event config modify`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
mail_from	string	Mail from
mail_server	string	Mail server (SMTP)
proxy_password	string	Password for HTTP/HTTPS proxy
proxy_url	string	HTTP/HTTPS proxy URL
proxy_user	string	User name for HTTP/HTTPS proxy
pubsub_enabled	boolean	Is Publish/Subscribe Messaging Enabled?

Example request

```
{
  "mail_from": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "mail_server": "<a href="
mailto:mail@mycompany.com">mail@mycompany.com</a>",
  "proxy_password": "password",
  "proxy_url": "https://proxyserver.mycompany.com",
  "proxy_user": "proxy_user",
  "pubsub_enabled": 1
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983123	The validation of the mail server provided failed
983136	The proxy URL cannot contain a username or password
983137	The proxy URL provided is invalid
983139	The IPv6 proxy URL provided is invalid
983140	The proxy URL provided contains an invalid scheme. Supported schemes are 'http' or 'https'

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

ems_config

Name	Type	Description
mail_from	string	Mail from
mail_server	string	Mail server (SMTP)
proxy_password	string	Password for HTTP/HTTPS proxy
proxy_url	string	HTTP/HTTPS proxy URL
proxy_user	string	User name for HTTP/HTTPS proxy
pubsub_enabled	boolean	Is Publish/Subscribe Messaging Enabled?

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage event destinations

Support EMS destinations endpoint overview

Overview

Manages the list of destinations. A destination is defined by a type and a place to which an event's information is transmitted.

Destination Types

An EMS destination is defined by a single type, which is one of the following:

- email
- syslog
- rest_api
- snmp

Email

The 'email' type allows you to define a mailbox where information about an observed event is sent by SMTP. Enter the address in the destination field in a valid format.

For example: administrator@mycompany.com

Syslog

The 'syslog' type allows you to specify a remote syslog server that can receive information about an observed event. Enter the hostname or IP address in the destination field.

For example: [syslog.mycompany.com](#), [192.168.1.1](#)

Rest API

The 'rest_api' type allows you to define a URL where information about an observed event is sent using the REST protocol. Enter the URL in the destination field.

The URL must contain a valid transmission schema which can be one of the following:

- http
 - https
- Using the 'https' schema, you can configure a client-side certificate for mutual authentication.
For example: <http://rest.mycompany.com>, <https://192.168.1.1>

SNMP

The 'snmp' type describes addresses where information about an observed event is sent using SNMP traps. The system defines a default instance of this type and it is restricted to read-only.

This type has the following limitations:

- Cannot create new destinations of type 'snmp'
- Cannot modify the default 'snmp' destination

SNMP trap host details need to be configured through one of the following:

Type	Command / API
CLI	'system snmp traphost'
ZAPI	'snmp-traphost-add' / 'snmp-traphost-delete'

Examples

Retrieving the list of active destinations

```
# API
GET /api/support/ems/destinations

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "name": "snmp-traphost",
      "_links": {
        "self": {
          "href": "/api/support/ems/destinations/snmp-traphost"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/support/ems/destinations"
    }
  }
}
```

Creating a new 'email' destination

```
# API
POST /support/ems/destinations

# JSON Body
{
  "name": "Technician_Email",
  "type": "email",
  "destination": "technician@mycompany.com",
  "filters": [
    { "name" : "critical-wafl" }
  ]
}

# Response
201 Created
```

Retrieve a collection of event destinations

GET /support/ems/destinations

Introduced In: 9.6

Retrieves a collection of event destinations.

Related ONTAP commands

- event notification destination show
- event notification show

Parameters

Name	Type	In	Required	Description
certificate.ca	string	query	False	Filter by certificate.ca
certificate.serial_number	string	query	False	Filter by certificate.serial_number
name	string	query	False	Filter by name
filters.name	string	query	False	Filter by filters.name
type	string	query	False	Filter by type

Name	Type	In	Required	Description
destination	string	query	False	Filter by destination
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "3",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "certificate": {
        "ca": "VeriSign",
        "serial_number": "1234567890"
      },
      "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
      "filters": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "name": "important-events"
        }
      ],
      "name": "Admin_Email",
      "type": "email"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

certificate

Certificate information is valid for the "rest_api" type.

Name	Type	Description
ca	string	Client certificate issuing CA
serial_number	string	Client certificate serial number

filters

Name	Type	Description
_links	_links	
name	string	

records

Name	Type	Description
_links	_links	
certificate	certificate	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[filters]	

Name	Type	Description
name	string	Destination name. Valid in POST.
type	string	Type of destination. Valid in POST.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create an event destination

POST /support/ems/destinations

Introduced In: 9.6

Creates an event destination.

Required properties

- `name` - String that uniquely identifies the destination.
- `type` - Type of destination that is to be created.
- `destination` - String that identifies the destination. The contents of this property changes depending on `type`.

Recommended optional properties

- `filters.name` - List of filter names that should direct to this destination.
- `certificate` - When specifying a rest api destination, a client certificate can be provided.

Related ONTAP commands

- `event notification destination create`
- `event notification create`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
certificate	certificate	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[filters]	
name	string	Destination name. Valid in POST.
type	string	Type of destination. Valid in POST.

Example request

```
{
  "certificate": {
    "ca": "VeriSign",
    "serial_number": "1234567890"
  },
  "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "filters": [
    {
      "name": "important-events"
    }
  ],
  "name": "Admin_Email",
  "type": "email"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "num_records": "3",
  "records": [
    {
      "certificate": {
        "ca": "VeriSign",
        "serial_number": "1234567890"
      },
      "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
      "filters": [
        {
          "name": "important-events"
        }
      ],
      "name": "Admin_Email",
      "type": "email"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983088	The destination name provided cannot be empty
983089	The destination name provided cannot contain spaces
983094	The destination name provided is invalid. The destination name must contain between 2 and 64 characters and start and end with an alphanumeric symbol or _(underscore). The allowed special characters are _(underscore) and -(hyphen)
983104	The syslog destination provided is invalid
983116	The number of notifications has reached maximum capacity
983117	The number of destinations has reached maximum capacity

Error Code	Description
983129	The rest-api destination provided must contain a valid scheme, such as http// or https//
983130	The rest-api destination provided contains an invalid URL
983131	The rest-api destination provided contains an invalid IPv6 URL
983144	The security certificate information provided is incomplete. Provide the certificate and serial number
983145	The rest-api destination provided has an 'http://' scheme. It is invalid to provide certificate information
983149	New SNMP destinations cannot be created
983151	A property provided cannot be configured on the requested destination type
983152	Default destinations cannot be modified or removed
983153	The security certificate provided does not exist
983154	The necessary private key is not installed on the system

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

certificate

Certificate information is valid for the "rest_api" type.

Name	Type	Description
ca	string	Client certificate issuing CA
serial_number	string	Client certificate serial number

filters

Name	Type	Description
name	string	

ems_destination

Name	Type	Description
certificate	certificate	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[filters]	
name	string	Destination name. Valid in POST.
type	string	Type of destination. Valid in POST.

records

Name	Type	Description
certificate	certificate	Certificate information is valid for the "rest_api" type.
destination	string	Event destination

Name	Type	Description
filters	array[filters]	
name	string	Destination name. Valid in POST.
type	string	Type of destination. Valid in POST.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage an event destination instance

Support EMS destinations name endpoint overview

Overview

Manages a specific instance of a destination. There are limits to the information that you can modify after a destination is created. For example, you cannot change a destination's type, but you can modify the underlying details of the type.

See the documentation for </support/ems/destinations> for details on the various properties in a destination.

Examples

Retrieving a specific destination instance

```
# API
GET /api/support/ems/destinations/snmp-traphost

# Response
200 OK

# JSON Body
{
  "name": "snmp-traphost",
  "type": "snmp",
  "destination": "",
  "filters": [
    {
      "name": "default-trap-events",
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/default-trap-events"
        }
      }
    }
  ],
  "_links": {
    "self": {
      "href": "/api/support/ems/destinations/snmp-traphost"
    }
  }
}
```

Updating an existing destination (change of email address)

```
# API
PATCH /api/support/ems/destinations/test-destination

# JSON Body
{
  "destination": "support@mycompany.com"
}

# Response
200 OK
```

Deleting an existing destination


```
# API
DELETE /api/support/ems/destinations/test-destination

# Response
200 OK
```

Delete an event destination

DELETE /support/ems/destinations/{name}

Introduced In: 9.6

Deletes an event destination.

Related ONTAP commands

- event notification destination delete
- event notification delete

Parameters

Name	Type	In	Required	Description
name	string	path	True	Destination name

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983152	Default destinations cannot be modified or removed

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve an event destination

GET /support/ems/destinations/{name}

Introduced In: 9.6

Retrieves event destinations.

Related ONTAP commands

- `event notification destination show`
- `event notification show`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Destination name
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
certificate	certificate	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[filters]	
name	string	Destination name. Valid in POST.
type	string	Type of destination. Valid in POST.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "certificate": {
    "ca": "VeriSign",
    "serial_number": "1234567890"
  },
  "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "filters": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "important-events"
    }
  ],
  "name": "Admin_Email",
  "type": "email"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

certificate

Certificate information is valid for the "rest_api" type.

Name	Type	Description
ca	string	Client certificate issuing CA
serial_number	string	Client certificate serial number

filters

Name	Type	Description
_links	_links	
name	string	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message

Name	Type	Description
target	string	The target parameter that caused the error.

Update an event destination

PATCH /support/ems/destinations/{name}

Introduced In: 9.6

Updates an event destination.

Recommended optional properties

- `filters.name` - New list of filters that should direct to this destination. The existing list is discarded.
- `certificate` - New certificate parameters when the destination type is `rest_api`.

Related ONTAP commands

- `event notification destination modify`
- `event notification modify`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Destination name

Request Body

Name	Type	Description
certificate	certificate	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[filters]	

Example request

```
{
  "certificate": {
    "ca": "VeriSign",
    "serial_number": "1234567890"
  },
  "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "filters": [
    {
      "name": "important-events"
    }
  ]
}
```

Response

Status: 200, Ok

Name	Type	Description
certificate	certificate	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[filters]	

Example response

```
{
  "certificate": {
    "ca": "VeriSign",
    "serial_number": "1234567890"
  },
  "destination": "<a href="
mailto:administrator@mycompany.com">administrator@mycompany.com</a>",
  "filters": [
    {
      "name": "important-events"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983088	The destination name provided cannot be empty
983089	The destination name provided cannot contain spaces
983094	The destination name provided is invalid. The destination name must contain between 2 and 64 characters and start and end with an alphanumeric symbol or _(underscore). The allowed special characters are _(underscore) and -(hyphen)
983104	The syslog destination provided is invalid
983116	The number of notifications has reached maximum capacity
983117	The number of destinations has reached maximum capacity
983129	The rest-api destination must contain a valid scheme, such as http// or https//
983130	The rest-api destination provided contains an invalid URL
983131	The rest-api destination provided contains an invalid IPv6 URL

Error Code	Description
983142	The security certificate provided does not exist
983143	The private security key provided does not exist
983144	The security certificate information provided is incomplete. Provide the certificate and serial number
983145	The rest-api destination provided has an 'http://' scheme. It is invalid to provide certificate information
983150	The type of an existing destination cannot be changed
983151	A property provided cannot be configured on the requested destination type
983152	Default destinations cannot be modified or removed

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

certificate

Certificate information is valid for the "rest_api" type.

Name	Type	Description
ca	string	Client certificate issuing CA
serial_number	string	Client certificate serial number

filters

Name	Type	Description
name	string	

ems_destination

Name	Type	Description
certificate	certificate	Certificate information is valid for the "rest_api" type.
destination	string	Event destination
filters	array[filters]	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve events

Support EMS events endpoint overview

Overview

Queries a live collection of observed events on the system.

Example

Querying for the latest event received by EMS

```
# API
GET /api/support/ems/events?fields=message.name&max_records=1

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "node": {
        "name": "node1",
        "uuid": "f087b8e3-99ac-11e8-b5a5-005056bb4ec7",
        "_links": {
          "self": {
            "href": "/api/cluster/nodes/f087b8e3-99ac-11e8-b5a5-005056bb4ec7"
          }
        }
      },
      "index": 661,
      "message": {
        "name": "raid.aggr.log.CP.count"
      },
      "_links": {
        "self": {
          "href": "/api/support/ems/events/node1/661"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/support/ems/events?fields=message.name&max_records=1"
    }
  }
}
```

Retrieve observed events

GET /support/ems/events

Introduced In: 9.6

Retrieves a collection of observed events.

Related ONTAP commands

- `event log show`

Parameters

Name	Type	In	Required	Description
log_message	string	query	False	Filter by log_message
parameters.name	string	query	False	Filter by parameters.name
parameters.value	string	query	False	Filter by parameters.value
message.severity	string	query	False	Filter by message.severity
message.name	string	query	False	Filter by message.name
source	string	query	False	Filter by source
index	integer	query	False	Filter by index
node.name	string	query	False	Filter by node.name
node.uuid	string	query	False	Filter by node.uuid
time	string	query	False	Filter by time
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. <ul style="list-style-type: none">• Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "3",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "index": "1",
      "log_message": "string",
      "message": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "callhome.spares.low",
        "severity": "emergency"
      },
      "node": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "node1",
        "uuid": "1cd8a442-86d1-11e0-ae1c-123478563412"
      },
      "parameters": [
        {
          "name": "numOps",
          "value": "123"
        }
      ],
      "source": "string",
```



```
    "time": "string"
  }
]
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

message

Name	Type	Description
_links	_links	
name	string	Message name of the event. Returned by default.
severity	string	Severity of the event. Returned by default.

node

Name	Type	Description
_links	_links	
name	string	
uuid	string	

parameters

Name	Type	Description
name	string	Name of parameter
value	string	Value of parameter

records

Name	Type	Description
_links	_links	
index	integer	Index of the event. Returned by default.
log_message	string	A formatted text string populated with parameter details. Returned by default.
message	message	
node	node	
parameters	array[parameters]	A list of parameters provided with the EMS event.
source	string	Source
time	string	Timestamp of the event. Returned by default.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create or retrieve event filters

Support EMS filters endpoint overview

Overview

Manages the list of available filters. A filter is a named collection of rules that enable the system to identify events that require additional handling. A filter is linked with a destination to which the system sends specific events.

When EMS processes an event, each filter is evaluated for a match. More than one filter can handle a single event.

Filter rule position

A filter's rules are evaluated sequentially, according to their position index. When a rule is added or modified, the position can be set to customize the filter's logic. If no position is specified, a new rule is appended to the end of the list.

Filter rule types

A filter rule can be one of two types: 'include' or 'exclude'.

If an event matches the criteria of the rule, the type dictates whether it should be forwarded to the destination or ignored.

Filter rule matching criteria

A valid filter rule must contain at least one set of criteria.

Name pattern

A name pattern is matched against an event's name. Multiple characters can be matched using the wildcard character '*'.

Severity

The severity pattern is matched against an event's severity. Multiple severities can be specified in a comma separated list. A single wildcard '*' will match all severities. Valid values are:

- emergency
- alert
- error
- notice
- informational
- debug

SNMP trap type

The SNMP trap type pattern is matched against an event's trap type. Multiple trap types can be specified in a comma separated list. A single wildcard '*' matches all trap types. Valid values are:

- standard
- built_in

- severity_based

Examples

Retrieving a list of filters whose names contain a hyphen

```
# API
GET /api/support/ems/filters?name=*-*
```

```
# Response
200 OK
```

```
# JSON Body
{
  "records": [
    {
      "name": "default-trap-events",
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/default-trap-events"
        }
      }
    },
    {
      "name": "important-events",
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/important-events"
        }
      }
    },
    {
      "name": "no-info-debug-events",
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/no-info-debug-events"
        }
      }
    }
  ],
  "num_records": 3,
  "_links": {
    "self": {
      "href": "/api/support/ems/filters?name=*-*"
    }
  }
}
```

Creating a new filter using various matching criteria

```
# API
POST /api/support/ems/filters

# JSON Body
{
  "name": "test-filter",
  "rules": [
    {
      "index": 1,
      "type": "include",
      "message_criteria": {
        "name_pattern": "LUN.*",
        "severities": "alert,error",
        "snmp_trap_types": "severity_based"
      }
    }
  ]
}

# Response
201 Created
```

Retrieve event filters

GET /support/ems/filters

Introduced In: 9.6

Retrieves a collection of event filters.

Related ONTAP commands

- event filter show

Parameters

Name	Type	In	Required	Description
rules.message_criteria.severities	string	query	False	Filter by rules.message_criteria.severities
rules.message_criteria.snmp_trap_types	string	query	False	Filter by rules.message_criteria.snmp_trap_types

Name	Type	In	Required	Description
rules.message_criteria.name_pattern	string	query	False	Filter by rules.message_criteria.name_pattern
rules.index	integer	query	False	Filter by rules.index
rules.type	string	query	False	Filter by rules.type
name	string	query	False	Filter by name
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "3",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "name": "snmp-traphost",
      "rules": [
        {
          "_links": {
            "self": {
              "href": "/api/resourcelink"
            }
          },
          "index": "1",
          "message_criteria": {
            "_links": {
              "related": {
                "href": "/api/resourcelink"
              }
            },
            "name_pattern": "callhome.*",
            "severities": "error,informational",
            "snmp_trap_types": "standard|built_in"
          },
          "type": "include"
        }
      ]
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

_links

Name	Type	Description
related	href	

message_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	_links	
name_pattern	string	Message name filter on which to match. Supports wildcards. Defaults to * if not specified.
severities	string	A comma-separated list of severities or a wildcard.
snmp_trap_types	string	A comma separated list of snmp_trap_types or a wildcard.

rules

Rule for an event filter

Name	Type	Description
_links	_links	

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

records

Name	Type	Description
_links	_links	
name	string	Filter name
rules	array[rules]	Array of event filter rules on which to match.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create an event filter

POST /support/ems/filters

Introduced In: 9.6

Creates an event filter.

Required properties

- `name` - String that uniquely identifies the filter.

Recommended optional properties

- `rules` - List of criteria which is used to match a filter with an event.

Related ONTAP commands

- `event filter create`

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
name	string	Filter name
rules	array[rules]	Array of event filter rules on which to match.

Example request

```
{
  "name": "snmp-traphost",
  "rules": [
    {
      "index": "1",
      "message_criteria": {
        "name_pattern": "callhome.*",
        "severities": "error,informational",
        "snmp_trap_types": "standard|built_in"
      },
      "type": "include"
    }
  ]
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "num_records": "3",
  "records": [
    {
      "name": "snmp-traphost",
      "rules": [
        {
          "index": "1",
          "message_criteria": {
            "name_pattern": "callhome.*",
            "severities": "error,informational",
            "snmp_trap_types": "standard|built_in"
          },
          "type": "include"
        }
      ]
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983088	The filter name provided is empty
983089	The filter name provided cannot contain spaces
983092	The index of the rule provided is outside the allowed range for the filter provided
983094	The filter name provided is invalid. The filter name must contain between 2 and 64 characters and start and end with an alphanumeric symbol or (underscore). The allowed special characters are (underscore) and -(hyphen)
983095	The rule index provided is invalid for the filter provided
983101	No event is matched by the rule provided
983113	Default filters cannot be modified or removed

Error Code	Description
983114	The maximum number of filters is reached
983115	The maximum number of filter rules is reached
983126	A rule requires at least one name_pattern, severities, or snmp_trap_types to be defined
983127	A property cannot contain a combination of the wildcard character and other values
983128	An invalid value is provided for the property 'snmp_trap_types'
983146	An invalid value is provided for the property 'severities'
983147	The severity provided are not supported
983155	The provided severities property does not match that of the name_pattern
983156	The provided snmp_trap_types property does not match that of the name_pattern
983157	The provided severities and snmp_trap_types do not match those of the name_pattern
983158	The name_pattern provided does not exist

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

message_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
name_pattern	string	Message name filter on which to match. Supports wildcards. Defaults to * if not specified.
severities	string	A comma-separated list of severities or a wildcard.
snmp_trap_types	string	A comma separated list of snmp_trap_types or a wildcard.

rules

Rule for an event filter

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

ems_filter

Name	Type	Description
name	string	Filter name

Name	Type	Description
rules	array[rules]	Array of event filter rules on which to match.

records

Name	Type	Description
name	string	Filter name
rules	array[rules]	Array of event filter rules on which to match.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage event filters

Support EMS filters name endpoint overview

Overview

Manages a specific filter instance.

See the documentation for </support/ems/filters> for details on the various properties.

Examples

Retrieving a specific filter instance

```
# API
GET /api/support/ems/filters/no-info-debug-events

# Response
200 OK

# JSON Body
{
  "name": "no-info-debug-events",
  "rules": [
    {
      "index": 1,
      "type": "include",
      "message_criteria": {
        "name_pattern": "*",
        "severities": "emergency,alert,error,notice",
        "snmp_trap_types": "*",
        "_links": {
          "related": {
            "href":
"/api/support/ems/messages?name=&severity=emergency,alert,error,notice&snmp_trap_type="
          }
        }
      },
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/no-info-debug-events/rules/1"
        }
      }
    },
    {
      "index": 2,
      "type": "exclude",
      "message_criteria": {
        "name_pattern": "*",
        "severities": "*",
        "snmp_trap_types": "*",
        "_links": {
          "related": {
            "href":
"/api/support/ems/messages?name=&severity=&snmp_trap_type="
          }
        }
      }
    }
  ]
}
```

```

    },
    "_links": {
      "self": {
        "href": "/api/support/ems/filters/no-info-debug-events/rules/2"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/support/ems/filters/no-info-debug-events"
    }
  }
}

```

Updating an existing filter with a new rule

```

# API
PATCH /api/support/ems/filters/test-filter

# JSON Body
{
  "rules": [
    {
      "type": "include",
      "message_criteria": {
        "name_pattern": "wafl.*",
        "severities": "error"
      }
    }
  ]
}

# Response
200 OK

```

Deleting an existing filter

```

# API
DELETE /api/support/ems/filters/test-filter

# Response
200 OK

```

Delete an event filter

DELETE /support/ems/filters/{name}

Introduced In: 9.6

Deletes an event filter.

Related ONTAP commands

- `event filter delete`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter name

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983113	Default filters cannot be modified or removed

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve an event filter

GET /support/ems/filters/{name}

Introduced In: 9.6

Retrieves an event filter.

Related ONTAP commands

- `event filter show`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter name
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
name	string	Filter name
rules	array[rules]	Array of event filter rules on which to match.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "name": "snmp-traphost",
  "rules": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "index": "1",
      "message_criteria": {
        "_links": {
          "related": {
            "href": "/api/resourcelink"
          }
        },
        "name_pattern": "callhome.*",
        "severities": "error,informational",
        "snmp_trap_types": "standard|built_in"
      },
      "type": "include"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

_links

Name	Type	Description
related	href	

message_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	_links	
name_pattern	string	Message name filter on which to match. Supports wildcards. Defaults to * if not specified.
severities	string	A comma-separated list of severities or a wildcard.
snmp_trap_types	string	A comma separated list of snmp_trap_types or a wildcard.

rules

Rule for an event filter

Name	Type	Description
_links	_links	
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.

Name	Type	Description
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update an event filter

PATCH /support/ems/filters/{name}

Introduced In: 9.6

Updates an event filter.

Recommended optional properties

- `new_name` - New string that uniquely identifies a filter.
- `rules` - New list of criteria used to match the filter with an event. The existing list is discarded.

Related ONTAP commands

- `event filter create`
- `event filter delete`
- `event filter rename`

- event filter rule add
- event filter rule delete
- event filter rule reorder

Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter name
new_name	string	query	False	New filter name for renames. Valid in PATCH.

Request Body

Name	Type	Description
name	string	Filter name
rules	array[rules]	Array of event filter rules on which to match.

Example request

```
{
  "name": "snmp-traphost",
  "rules": [
    {
      "index": "1",
      "message_criteria": {
        "name_pattern": "callhome.*",
        "severities": "error,informational",
        "snmp_trap_types": "standard|built_in"
      },
      "type": "include"
    }
  ]
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983088	The filter name provided is empty
983089	The filter name provided cannot contain spaces
983092	The index of the rule provided is outside the allowed range for the filter provided
983094	The filter name provided is invalid. The filter name must contain between 2 and 64 characters and start and end with an alphanumeric symbol or (underscore). The allowed special characters are (underscore) and -(hyphen)
983095	The rule index provided is invalid for the filter provided
983101	No event is matched by the rule provided
983113	Default filters cannot be modified or removed
983114	The maximum number of filters is reached
983115	The maximum number of filter rules is reached
983126	A rule requires at least one of name_pattern, severities or snmp-trap-types to be defined
983127	A property cannot contain a combination of the wildcard characters and other values
983128	An invalid value is provided for the property 'snmp_trap_types'
983146	An invalid value is provided for the property 'severities'
983147	The severity levels provided are not supported
983155	The provided severities property does not match that of the name_pattern
983156	The provided snmp_trap_types property does not match that of the name_pattern
983157	The provided severities and snmp_trap_types do not match those of the name_pattern
983158	The name_pattern provided does not exist

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

message_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
name_pattern	string	Message name filter on which to match. Supports wildcards. Defaults to * if not specified.
severities	string	A comma-separated list of severities or a wildcard.
snmp_trap_types	string	A comma separated list of snmp_trap_types or a wildcard.

rules

Rule for an event filter

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

ems_filter

Name	Type	Description
name	string	Filter name

Name	Type	Description
rules	array[rules]	Array of event filter rules on which to match.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create or retrieve filter rules

Support EMS filters name rules endpoint overview

Overview

Manages the list of rules associated with a specific filter. A filter contains a list of rules that are evaluated to determine whether an event matches the filter. When a rule matches an event, the filter is considered a match. See the documentation for </support/ems/filters> for details on the various properties in a rule.

Examples

Retrieving the collection of rules connected to a filter


```
# API
GET /api/support/ems/filters/no-info-debug-events/rules

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "index": 1,
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/no-info-debug-events/rules/1"
        }
      }
    },
    {
      "index": 2,
      "_links": {
        "self": {
          "href": "/api/support/ems/filters/no-info-debug-events/rules/2"
        }
      }
    }
  ],
  "num_records": 2,
  "_links": {
    "self": {
      "href": "/api/support/ems/filters/no-info-debug-events/rules"
    }
  }
}
```

Creating a new rule under an existing filter



The absence of a position index appends the rule to the end of the list.

```
# API
POST /api/support/ems/filters/test-filter/rules

# JSON Body
{
  "type": "include",
  "message_criteria": {
    "name_pattern": "wafl.*",
    "severities": "error,informational"
  }
}

# Response
201 Created
```

Retrieve event filter rules

GET /support/ems/filters/{name}/rules

Introduced In: 9.6

Retrieves event filter rules.

Related ONTAP commands

- `event filter show`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter Name
message_criteria.severities	string	query	False	Filter by message_criteria.severities
message_criteria.snmp_trap_types	string	query	False	Filter by message_criteria.snmp_trap_types
message_criteria.name_pattern	string	query	False	Filter by message_criteria.name_pattern
index	integer	query	False	Filter by index

Name	Type	In	Required	Description
type	string	query	False	Filter by type
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	

Name	Type	Description
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "3",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "index": "1",
      "message_criteria": {
        "_links": {
          "related": {
            "href": "/api/resourcelink"
          }
        },
        "name_pattern": "callhome.*",
        "severities": "error,informational",
        "snmp_trap_types": "standard|built_in"
      },
      "type": "include"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

_links

Name	Type	Description
related	href	

message_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	_links	
name_pattern	string	Message name filter on which to match. Supports wildcards. Defaults to * if not specified.
severities	string	A comma-separated list of severities or a wildcard.
snmp_trap_types	string	A comma separated list of snmp_trap_types or a wildcard.

records

Rule for an event filter

Name	Type	Description
_links	_links	

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create an event filter rule

POST /support/ems/filters/{name}/rules

Introduced In: 9.6

Creates an event filter rule.

Required properties

- `message_criteria` - Criteria on which a rule is to match an event.

Recommended optional properties

- `index` - One-based position index of the new rule.

Related ONTAP commands

- `event filter rule add`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter name
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

Example request

```
{
  "index": "1",
  "message_criteria": {
    "name_pattern": "callhome.*",
    "severities": "error,informational",
    "snmp_trap_types": "standard|built_in"
  },
  "type": "include"
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "num_records": "3",
  "records": [
    {
      "index": "1",
      "message_criteria": {
        "name_pattern": "callhome.*",
        "severities": "error,informational",
        "snmp_trap_types": "standard|built_in"
      },
      "type": "include"
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983092	The index of the rule provided is outside the allowed range for the filter provided
983095	The rule index provided is invalid for the filter provided
983113	Default filters cannot be modified or removed
983115	The maximum number of filter rules is reached
983126	A rule requires at least one name_pattern, severities, or snmp_trap_types to be defined
983127	A property cannot contain a combination of the wildcard characters and other values
983128	An invalid value is provided for the property 'snmp_trap_types'
983146	An invalid value is provided for the property 'severities'
983147	The severity levels provided are not supported
983155	The provided severities property does not match that of the name_pattern
983156	The provided snmp_trap_types property does not match that of the name_pattern
983157	The provided severities and snmp_trap_types do not match those of the name_pattern
983158	The name_pattern provided does not exist

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

message_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
name_pattern	string	Message name filter on which to match. Supports wildcards. Defaults to * if not specified.
severities	string	A comma-separated list of severities or a wildcard.
snmp_trap_types	string	A comma separated list of snmp_trap_types or a wildcard.

ems_filter_rule

Rule for an event filter

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

records

Rule for an event filter

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage event filter rules

Support EMS filters name rules index endpoint overview

Overview

Manages a specific instance of a rule within a filter.

See the documentation for </support/ems/filters> for details on the various properties in a rule.

Examples

Retrieving a single instance of a rule

```
# API
GET /api/support/ems/filters/no-info-debug-events/rules/1

# Response
200 OK

# JSON Body
{
  "name": "no-info-debug-events",
  "index": 1,
  "type": "include",
  "message_criteria": {
    "name_pattern": "*",
    "severities": "emergency,alert,error,notice",
    "snmp_trap_types": "*",
    "_links": {
      "self": {
        "href":
"/api/support/ems/messages?name=&severity=emergency,alert,error,notice&snmp_trap_type="
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/support/ems/filters/no-info-debug-events/rules/1"
    }
  }
}
```

Updating an existing rule to use severity emergency

```
# API
PATCH /api/support/ems/filters/test-filter/rules/1

# JSON Body
{
  "message_criteria": {
    "severities": "emergency"
  }
}

# Response
200 OK
```

Deleting a rule from an existing filter

```
# API
DELETE /api/support/ems/filters/test-filter/rules/1

# Response
200 OK
```

Delete an event filter rule

DELETE /support/ems/filters/{name}/rules/{index}

Introduced In: 9.6

Deletes an event filter rule.

Related ONTAP commands

- event filter rule delete

Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter name
index	string	path	True	Filter index

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983091	A default rule cannot be removed
983092	The index of the rule provided is outside the allowed range for the filter provided
983095	The rule index provided is invalid for the filter provided
983110	There are no user defined rules in the filter provided
983113	Default filters cannot be modified or removed

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve an event filter rule

GET /support/ems/filters/{name}/rules/{index}

Introduced In: 9.6

Retrieves an event filter rule.

Related ONTAP commands

- `event filter show`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter name
index	string	path	True	Filter index
fields	array[string]	query	False	Specify the fields to return.

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "index": "1",
  "message_criteria": {
    "_links": {
      "related": {
        "href": "/api/resourcelink"
      }
    },
    "name_pattern": "callhome.*",
    "severities": "error,informational",
    "snmp_trap_types": "standard|built_in"
  },
  "type": "include"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

_links

Name	Type	Description
related	href	

message_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
_links	_links	
name_pattern	string	Message name filter on which to match. Supports wildcards. Defaults to * if not specified.
severities	string	A comma-separated list of severities or a wildcard.
snmp_trap_types	string	A comma separated list of snmp_trap_types or a wildcard.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array [error_arguments]	Message arguments

Name	Type	Description
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update an event filter rule

PATCH /support/ems/filters/{name}/rules/{index}

Introduced In: 9.6

Updates an event filter rule.

Recommended optional properties

- `message_criteria` - New criteria on which a rule is to match an event.

Related ONTAP commands

- `event filter rule add`
- `event filter rule delete`

Parameters

Name	Type	In	Required	Description
name	string	path	True	Filter name
index	string	path	True	Filter index
new_index	integer	query	False	New position for the filter rule index

Request Body

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.

Name	Type	Description
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

Example request

```
{
  "index": "1",
  "message_criteria": {
    "name_pattern": "callhome.*",
    "severities": "error,informational",
    "snmp_trap_types": "standard|built_in"
  },
  "type": "include"
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
983092	The index of the rule provided is outside the allowed range for the filter provided
983095	The rule index provided is invalid for the filter provided
983113	Default filters cannot be modified or removed
983126	A rule requires at least one name_pattern, severities, or snmp_trap_types to be defined
983127	A property cannot contain a combination of the wildcard characters and other values.
983128	An invalid value is provided for the property 'snmp_trap_types'

Error Code	Description
983146	An invalid value is provided for the property 'severities'
983147	The severity levels provided are not supported
983155	The provided severities property does not match that of the name_pattern
983156	The provided snmp_trap_types property does not match that of the name_pattern
983157	The provided severities and snmp_trap_types do not match those of the name_pattern
983158	The name_pattern provided does not exist

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

message_criteria

Matching message definitions for the filter. A property must be specified.

Name	Type	Description
name_pattern	string	Message name filter on which to match. Supports wildcards. Defaults to * if not specified.
severities	string	A comma-separated list of severities or a wildcard.
snmp_trap_types	string	A comma separated list of snmp_trap_types or a wildcard.

ems_filter_rule

Rule for an event filter

Name	Type	Description
index	integer	Rule index. Rules are evaluated in ascending order. If a rule's index order is not specified during creation, the rule is appended to the end of the list.
message_criteria	message_criteria	Matching message definitions for the filter. A property must be specified.
type	string	Rule type

error_arguments

Name	Type	Description
code	string	Argument code

Name	Type	Description
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage EMS event catalog

Support EMS messages endpoint overview

Overview

Allows access to the EMS event catalog. The catalog contains a list of all events supported by the system and their corresponding descriptions, the reason for an event occurrence, and how to correct issues related to the event.

Example

Querying for the first event that has a message name beginning with 'C'


```
# API
GET /api/support/ems/messages?fields=name&max_records=1&name=C*

# Response
200 OK

# JSON Body
{
  "records": [
    {
      "name": "CR.Data.File.Inaccessible",
      "_links": {
        "self": {
          "href": "/api/support/ems/messages/CR.Data.File.Inaccessible"
        }
      }
    }
  ],
  "num_records": 1,
  "_links": {
    "self": {
      "href": "/api/support/ems/messages?fields=name&max_records=1&name=C*"
    }
  }
}
```

Retrieve event catalog definitions

GET /support/ems/messages

Introduced In: 9.6

Retrieves the event catalog definitions.

Related ONTAP commands

- event catalog show

Parameters

Name	Type	In	Required	Description
severity	string	query	False	Filter by severity
name	string	query	False	Filter by name

Name	Type	In	Required	Description
deprecated	boolean	query	False	Filter by deprecated
corrective_action	string	query	False	Filter by corrective_action
snmp_trap_type	string	query	False	Filter by snmp_trap_type
description	string	query	False	Filter by description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records
records	array[records]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "num_records": "3",
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "corrective_action": "string",
      "deprecated": 1,
      "description": "string",
      "name": "callhome.spares.low",
      "severity": "error",
      "snmp_trap_type": "standard"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

records

Name	Type	Description
_links	_links	
corrective_action	string	Corrective action
deprecated	boolean	Is deprecated?
description	string	Description
name	string	Name of the event.
severity	string	Severity
snmp_trap_type	string	SNMP trap type

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage cluster-wide SNMP configuration

Support SNMP endpoint overview

Overview

Cluster wide SNMP configuration. You can configure or retrieve the following SNMP parameters using this endpoint:

- enable or disable SNMP
- enable or disable SNMP authentication traps

Examples

Disables SNMP protocol in the cluster.

```
# The API:
PATCH "/api/support/snmp"

# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d '{"enabled":"false"}'

# The response
200 OK
{}
```

Enables SNMP authentication traps in the cluster.

```
# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d
'{"auth_traps_enabled":"true"}'

# The response
200 OK
{}
```

Enables SNMP protocol and SNMP authentication traps in the cluster.

```
# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-ip>/api/support/snmp" -d '{"enabled":"true",
"auth_traps_enabled":"true"}'

# The response
200 OK
{}
```

Retrieve the cluster-wide SNMP configuration

GET /support/snmp

Introduced In: 9.7

Retrieves the cluster wide SNMP configuration.

Related ONTAP commands

- `options snmp.enable`
- `system snmp show`

Learn more

- [DOC /support/snmp](#)

Parameters

Name	Type	In	Required	Description
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is true for GET calls. When set to false, only the number of records is returned. • Default value: 1
return_timeout	integer	query	False	The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached. • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	self_link	
auth_traps_enabled	boolean	Specifies whether to enable or disable SNMP authentication traps.
enabled	boolean	Specifies whether to enable or disable SNMP.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "auth_traps_enabled": 1,
  "enabled": 1
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

Name	Type	Description
self	href	

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the cluster-wide SNMP configuration

PATCH /support/snmp

Introduced In: 9.7

Updates the cluster wide SNMP configuration, such as enabling or disabling SNMP and enabling or disabling authentication traps.

Related ONTAP commands

- `options snmp.enable`
- `system snmp authtrap`

Learn more

- [DOC /support/snmp](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. <ul style="list-style-type: none">• Default value:

Request Body

Name	Type	Description
auth_traps_enabled	boolean	Specifies whether to enable or disable SNMP authentication traps.
enabled	boolean	Specifies whether to enable or disable SNMP.

Example request

```
{
  "auth_traps_enabled": 1,
  "enabled": 1
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
1179738	Cluster identification not available for modification since a cluster has not been created.

Error Code	Description
1179739	Cluster identification modification is not available in maintenance mode or other boot modes.

Definitions

See Definitions

href

Name	Type	Description
href	string	

self_link

snmp

Cluster-wide SNMP configuration.

Name	Type	Description
auth_traps_enabled	boolean	Specifies whether to enable or disable SNMP authentication traps.
enabled	boolean	Specifies whether to enable or disable SNMP.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create or retrieve SNMP traphosts

Support SNMP traphosts endpoint overview

Overview

This SNMP endpoint is used to retrieve, configure and reconfigure SNMP traphosts. An SNMP traphost is a host machine that is configured to receive SNMP traps from ONTAP. An SNMP traphost can be an SNMPv1/SNMPv2c traphost or an SNMPv3 traphost. ONTAP sends SNMPv1/SNMPv2c traps to an SNMPv1/SNMPv2c traphost over the SNMPv1/SNMPv2c protocol. Similarly, it sends SNMPv3 traps to an SNMPv3 traphost over the SNMPv3 protocol.

Examples

Retrieves a list of SNMP traphosts in the cluster, sorted by IP address

```
# The API:
GET "/api/support/snmp/traphosts"

# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/traphosts?order_by=ip_address"

# The response:
{
  "records": [
    {
      "host": "scspr0651011001.gdl.englab.netapp.com",
      "ip_address": "10.235.36.62",
      "_links": {
        "self": {
          "href":
"/api/support/snmp/traphosts/scspr0651011001.gdl.englab.netapp.com"
        }
      }
    },
    {
      "host": "scspr0647726006.gdl.englab.netapp.com",
      "ip_address": "10.235.36.139",
      "_links": {
        "self": {
          "href":
"/api/support/snmp/traphosts/scspr0647726006.gdl.englab.netapp.com"
        }
      }
    },
    {
      "host": "scspr0651055012.gdl.englab.netapp.com",
      "ip_address": "10.237.245.106",
      "_links": {
```

```
    "self": {
      "href":
"/api/support/snmp/traphosts/scspr0651055012.gdl.englab.netapp.com"
    }
  }
],
"num_records": 3,
"_links": {
  "self": {
    "href": "/api/support/snmp/traphosts?order_by=ip_address"
  }
}
```

Retrieves a list of SNMP traphosts in the cluster for which SNMP user is 'snmpv3user1'

```
# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/traphosts?name=snmpv3user1"

# The response:
{
  "records": [
    {
      "host": "scspr0653246020.gdl.englab.netapp.com",
      "user": {
        "name": "snmpv3user1",
        "_links": {
          "self": {
            "href":
"/api/support/snmp/users/800003150558b57e8dbd9ce9119d82005056a7b4e5/snmpv3
user1"
          }
        }
      },
      "_links": {
        "self": {
          "href":
"/api/support/snmp/traphosts/scspr0653246020.gdl.englab.netapp.com"
        }
      }
    },
    {
      "num_records": 1,
      "_links": {
        "self": {
          "href": "/api/support/snmp/traphosts?user.name=snmpv3user1"
        }
      }
    }
  ]
}
```

Adds an SNMPv1 or SNMPv2c traphost


```
# The API:
POST "/api/support/snmp/traphosts"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/traphosts" -d
'{"host":"10.140.128.48"}'

# The response:
201 Created
{}
```

Adds an SNMPv3 traphost

```
# The API:
POST "/api/support/snmp/traphosts"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/traphosts" -d '{"host":"fd-
geoycpi-uno.gycpi.b.yahoodns.net", "user":{"name":"snmpv3user1"}}'

# The response:
201 Created
{}
```

Retrieve SNMP traphosts

GET /support/snmp/traphosts

Introduced In: 9.7

Retrieves the list of SNMP traphosts along with the SNMP users configured for those traphosts.

Related ONTAP commands

- `system snmp traphost show`

Learn more

- [DOC /support/snmp/traphosts](#)

Parameters

Name	Type	In	Required	Description
host	string	query	False	Filter by host
user.name	string	query	False	Filter by user.name
ip_address	string	query	False	Filter by ip_address
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records.
records	array[snmp_traphost]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "host": "traphost.example.com",
      "ip_address": "10.10.10.7",
      "user": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "snmpv3user3"
      }
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

user

Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

Name	Type	Description
_links	_links	

Name	Type	Description
name	string	<p>Optional SNMPv1/SNMPv2c or SNMPv3 user name. For an SNMPv3 traphost, this object refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this object refers to an SNMP community. For an SNMPv3 traphost, this object is mandatory and refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, ONTAP automatically uses "public", if the same is configured, or any other configured community as user. So, for an SNMPv1 or SNMPv2c traphost, this property should not be provided in the "POST" method. However, the configured community for the SNMPv1/SNMPv2c traphost is returned by the "GET" method.</p> <ul style="list-style-type: none"> • example: snmpv3user3 • Introduced in: 9.7

snmp_traphost

SNMP manager or host machine that receives SNMP traps from ONTAP.

Name	Type	Description
_links	_links	
host	string	Fully qualified domain name (FQDN), IPv4 address or IPv6 address of SNMP traphost.
ip_address	string	IPv4 or IPv6 address
user	user	Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create SNMP traphosts

POST /support/snmp/traphosts

Introduced In: 9.7

Creates SNMP traphosts. While adding an SNMPv3 traphost, an SNMPv3 user configured in ONTAP must be specified. ONTAP uses this user's credentials to authenticate and/or encrypt traps sent to this SNMPv3 traphost. While adding an SNMPv1/SNMPv2c traphost, SNMPv1/SNMPv2c user or community need not be specified.

Required properties

- `host` - Fully Qualified Domain Name (FQDN), IPv4 address or IPv6 address of SNMP traphost.

Recommended optional properties

- If `host` refers to an SNMPv3 traphost, the following field is required:
 - `user` - SNMPv3 or User-based Security Model (USM) user.
- For an SNMPv1/SNMPv2c traphost, ONTAP automatically uses 'public' if 'public' is configured or no community is configured. Otherwise, ONTAP uses the first configured community.

Related ONTAP commands

- `system snmp traphost add`

Learn more

- [DOC /support/snmp/traphosts](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Request Body

Name	Type	Description
host	string	Fully qualified domain name (FQDN), IPv4 address or IPv6 address of SNMP traphost.
user	user	Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

Example request

```
{
  "host": "traphost.example.com",
  "user": {
    "name": "snmpv3user3"
  }
}
```

Response

Status: 201, Created

Name	Type	Description
num_records	integer	Number of records.
records	array[snmp_traphost]	

Example response

```
{
  "records": [
    {
      "host": "traphost.example.com",
      "user": {
        "name": "snmpv3user3"
      }
    }
  ]
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
8847365	Unknown host.
9043969	Duplicate traphost entry.
9043980	IPv6 support is disabled.
9043991	Not a USM user.
9043993	Current cluster version does not support SNMPv3 traps.
9044001	Failed to create SNMPv1 traphost.
9044002	Failed to create SNMPv3 traphost.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

user

Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

Name	Type	Description
name	string	<p>Optional SNMPv1/SNMPv2c or SNMPv3 user name. For an SNMPv3 traphost, this object refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this object refers to an SNMP community. For an SNMPv3 traphost, this object is mandatory and refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, ONTAP automatically uses "public", if the same is configured, or any other configured community as user. So, for an SNMPv1 or SNMPv2c traphost, this property should not be provided in the "POST" method. However, the configured community for the SNMPv1/SNMPv2c traphost is returned by the "GET" method.</p> <ul style="list-style-type: none">• example: snmpv3user3• Introduced in: 9.7

snmp_traphost

SNMP manager or host machine that receives SNMP traps from ONTAP.

Name	Type	Description
host	string	Fully qualified domain name (FQDN), IPv4 address or IPv6 address of SNMP traphost.
user	user	Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage individual SNMP traphosts

Support SNMP traphosts host endpoint overview

Overview

Defines, retrieves or deletes an individual SNMP traphost.

Examples

Retrieves an individual traphost in the cluster

```

# The API:
GET "/api/support/snmp/traphosts/{host}"

# The call
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/traphosts/10.235.36.62"

# The response
{
  "host": "scspr0651011001.gdl.englab.netapp.com",
  "ip_address": "10.235.36.62",
  "user": {
    "name": "public",
    "_links": {
      "self": {
        "href":
"/api/support/snmp/users/800003150558b57e8dbd9ce9119d82005056a7b4e5/public
"
      }
    }
  },
  "_links": {
    "self": {
      "href": "/api/support/snmp/traphosts/10.235.36.62"
    }
  }
}

```

Deletes an individual traphost in the cluster

```

# The API:
DELETE "/api/support/snmp/traphosts/{host}"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
DELETE "https://<mgmt-
ip>/api/support/snmp/traphosts/3ffe:ffff:100:f102::1"

# The response:
200 OK
{}

```

Delete an SNMP traphost

DELETE /support/snmp/traphosts/{host}

Introduced In: 9.7

Deletes an SNMP traphost.

Learn more

- [DOC /support/snmp/traphosts/{host}](#)

Parameters

Name	Type	In	Required	Description
host	string	path	True	Fully Qualified Domain Name (FQDN), IPv4 address or IPv6 address of SNMP traphost.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
9043970	Traphost entry does not exist.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve SNMP traphost details

GET /support/snmp/traphosts/{host}

Introduced In: 9.7

Retrieves the details of an SNMP traphost along with the SNMP user configured for that traphost.

Learn more

- [DOC /support/snmp/traphosts/{host}](#)

Parameters

Name	Type	In	Required	Description
host	string	path	True	Fully Qualified Domain Name (FQDN), IPv4 address or IPv6 address of SNMP traphost.
fields	array[string]	query	False	Specify the fields to return.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
host	string	Fully qualified domain name (FQDN), IPv4 address or IPv6 address of SNMP traphost.
ip_address	string	IPv4 or IPv6 address
user	user	Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "host": "traphost.example.com",
  "ip_address": "10.10.10.7",
  "user": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "snmpv3user3"
  }
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

user

Optional SNMP user parameter. For an SNMPv3 traphost, this property refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this property refers to an SNMP community.

Name	Type	Description
_links	_links	
name	string	<p>Optional SNMPv1/SNMPv2c or SNMPv3 user name. For an SNMPv3 traphost, this object refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, this object refers to an SNMP community. For an SNMPv3 traphost, this object is mandatory and refers to an SNMPv3 or User-based Security Model (USM) user. For an SNMPv1 or SNMPv2c traphost, ONTAP automatically uses "public", if the same is configured, or any other configured community as user. So, for an SNMPv1 or SNMPv2c traphost, this property should not be provided in the "POST" method. However, the configured community for the SNMPv1/SNMPv2c traphost is returned by the "GET" method.</p> <ul style="list-style-type: none">• example: snmpv3user3• Introduced in: 9.7

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve and create SNMP users

Support SNMP users endpoint overview

Overview

Defines, retrieves and adds new SNMP users. An SNMP user can be an SNMPv1/SNMPv2c user or an SNMPv3 user. SNMPv1/SNMPv2c user is also called a "community" user. Use a "community" user to query ONTAP SNMP server over SNMPv1/SNMPv2c protocol. An SNMPv3 user, also called a User-based Security Model (USM) user, can be a local SNMPv3 user or a remote SNMPv3 user. Use a local SNMPv3 user to query ONTAP SNMP server over SNMPv3 and/or to send SNMPv3 traps. A remote SNMPv3 user is configured in ONTAP as well as on a remote switch. ONTAP SNMP client uses a remote SNMPv3 user to query the switch over SNMPv3.

Examples

Retrieves a list of SNMP users in the cluster, sorted by SVM UUID

```
# The API:
GET "/api/support/snmp/users"

# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/users?order_by=owner.uuid"

# The response:
{
  "records": [
```

```

{
  "engine_id": "80000315052d7d2c4410b8e911bc9f005056bb942e",
  "name": "snmpv1user1",
  "owner": {
    "uuid": "442c7d2d-b810-11e9-bc9f-005056bb942e"
  },
  "_links": {
    "self": {
      "href":
"/api/support/snmp/users/80000315052d7d2c4410b8e911bc9f005056bb942e/snmpv1
user1"
    }
  }
},
{
  "engine_id": "80000315058e02057c0fb8e911bc9f005056bb942e",
  "name": "snmpv3user1",
  "owner": {
    "uuid": "7c05028e-b80f-11e9-bc9f-005056bb942e"
  },
  "_links": {
    "self": {
      "href":
"/api/support/snmp/users/80000315058e02057c0fb8e911bc9f005056bb942e/snmpv3
user1"
    }
  }
}
],
"num_records": 2,
"_links": {
  "self": {
    "href": "/api/support/snmp/users?order_by=owner.uuid"
  }
}
}

```

Creates an SNMPv1 or SNMPv2c community user

```
# The API:
POST "/api/support/snmp/users"

# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/users" -d
'{"owner":{"uuid":"02c9e252-41be-11e9-81d5-00a0986138f7",
"name":"cluster1"}, "name":"public", "authentication_method":"community"}'

# The response:
201 Created
{}
```

Creates a local SNMPv3 user

```
# The API:
POST "/api/support/snmp/users"

# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/users" -d
'{"owner":{"uuid":"02c9e252-41be-11e9-81d5-00a0986138f7",
"name":"cluster1"}, "name":"snmpv3u1", "authentication_method":"usm",
"snmpv3":{"authentication_protocol":"sha",
"authentication_password":"jelly22fi$h", "privacy_protocol":"aes",
"privacy_password":"allBlack$"}}'

# The response:
201 Created
{}
```

Creates a remote SNMPv3 user

```
# The API:
POST "/api/support/snmp/users"

# The call
curl -H "accept: application/json" -H "Content-Type: application/json" -X
POST "https://<mgmt-ip>/api/support/snmp/users" -d
'{"owner":{"uuid":"02c9e252-41be-11e9-81d5-00a0986138f7",
"name":"cluster1"},
"engine_id":"8000031505b67667a26975e9118a480050568e6f74",
"name":"switchuser1", "authentication_method":"usm",
"switch_address":"fdf8:f53b:82e4::53",
"snmpv3":{"authentication_protocol":"md5",
"authentication_password":"TfhIeliw601lS.Rw$3pm"}}'
```

Retrieve SNMP users on a cluster

GET /support/snmp/users

Introduced In: 9.7

Retrieves the list of SNMP users on the cluster.

Related ONTAP commands

- security snmpusers
- security login show -application snmp

Learn more

- [DOC /support/snmp/users](#)

Parameters

Name	Type	In	Required	Description
switch_address	string	query	False	Filter by switch_address
snmpv3.privacy_protocol	string	query	False	Filter by snmpv3.privacy_protocol

Name	Type	In	Required	Description
snmpv3.authentication_protocol	string	query	False	Filter by snmpv3.authentication_protocol
authentication_method	string	query	False	Filter by authentication_method
scope	string	query	False	Filter by scope
name	string	query	False	Filter by name
owner.uuid	string	query	False	Filter by owner.uuid
owner.name	string	query	False	Filter by owner.name
comment	string	query	False	Filter by comment
engine_id	string	query	False	Filter by engine_id
fields	array[string]	query	False	Specify the fields to return.
max_records	integer	query	False	Limit the number of records returned.
return_records	boolean	query	False	<p>The default is true for GET calls. When set to false, only the number of records is returned.</p> <ul style="list-style-type: none"> • Default value: 1

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0
order_by	array[string]	query	False	Order results by specified fields and optional [asc

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
num_records	integer	Number of records.
records	array[snmp_user]	

Example response

```
{
  "_links": {
    "next": {
      "href": "/api/resourcelink"
    },
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "records": [
    {
      "_links": {
        "self": {
          "href": "/api/resourcelink"
        }
      },
      "authentication_method": "usm",
      "comment": "This is a comment.",
      "engine_id": "80000315055415ab26d4aae811ac4d005056bb792e",
      "name": "snmpv3user2",
      "owner": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "name": "svm1",
        "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
      },
      "scope": "svm",
      "snmpv3": {
        "_links": {
          "self": {
            "href": "/api/resourcelink"
          }
        },
        "authentication_protocol": "sha2_256",
        "privacy_protocol": "aes128"
      },
      "switch_address": "10.23.34.45"
    }
  ]
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
next	href	
self	href	

_links

Name	Type	Description
self	href	

owner

Optional name and UUID of owning Storage Virtual Machine (SVM).

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snmpv3

Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
_links	_links	
authentication_protocol	string	Authentication protocol.
privacy_protocol	string	Privacy protocol.

snmp_user

An SNMP user can be an SNMPv1/SNMPv2c user or an SNMPv3 user. SNMPv1/SNMPv2c user is also called a "community" user. An SNMPv3 user, also called a User-based Security Model (USM) user, can be a local SNMPv3 user or a remote SNMPv3 user. A local SNMPv3 user can be used for querying ONTAP SNMP server over SNMPv3 and/or for sending SNMPv3 traps. The local SNMPv3 user used for

sending SNMPv3 traps must be configured with the same authentication and privacy credentials on the traphost receiver as well. A remote SNMPv3 user is also configured on a remote switch and used by ONTAP SNMP client functionality to query the remote switch over SNMPv3. An SNMP user is scoped to its owning Storage Virtual Machine (SVM). Owning SVM could be a data SVM or the administrative SVM.

Name	Type	Description
_links	_links	
authentication_method	string	Optional authentication method.
comment	string	Optional comment text.
engine_id	string	Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier.
name	string	SNMP user name.
owner	owner	Optional name and UUID of owning Storage Virtual Machine (SVM).
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
snmpv3	snmpv3	Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
switch_address	string	<p>Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.</p> <ul style="list-style-type: none"> • example: 10.23.34.45 • Introduced in: 9.7 • readCreate: 1

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Create a cluster-scoped or an SVM-scoped SNMP user

POST /support/snmp/users

Introduced In: 9.7

Creates either a cluster-scoped or an SVM-scoped SNMP user. This user can be an SNMPv1 or SNMPv2c community user or an SNMPv3 user. An SNMPv3 user can be a local SNMPv3 user or a remote SNMPv3 user.

Required properties

- `owner` - Name and UUID of owning SVM.
- `engine_id` - Engine ID of owning SVM or remote switch.
- `name` - SNMP user name
- `authentication_method` - Authentication method

Optional properties

- `switch_address` - Optional remote switch address
- `snmpv3` - SNMPv3-specific credentials
- `comment` - Comment text

Default property values

- `snmpv3.authentication_protocol` - none
- `snmpv3.privacy_protocol` - none

Related ONTAP commands

- `security login create`
- `system snmp community add`

Learn more

- [DOC /support/snmp/users](#)

Parameters

Name	Type	In	Required	Description
return_records	boolean	query	False	The default is false. If set to true, the records are returned. • Default value:

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
authentication_method	string	Optional authentication method.
comment	string	Optional comment text.

Name	Type	Description
engine_id	string	Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier.
name	string	SNMP user name.
owner	owner	Optional name and UUID of owning Storage Virtual Machine (SVM).
snmpv3	snmpv3	Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
switch_address	string	<p>Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.</p> <ul style="list-style-type: none"> • example: 10.23.34.45 • Introduced in: 9.7 • readCreate: 1

Example request

```
{
  "authentication_method": "usm",
  "comment": "This is a comment.",
  "engine_id": "80000315055415ab26d4aae811ac4d005056bb792e",
  "name": "snmpv3user2",
  "owner": {
    "name": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "snmpv3": {
    "authentication_password": "humTdumt*@t0nAwall",
    "authentication_protocol": "sha2_256",
    "privacy_password": "p@**GOandCLCt*200",
    "privacy_protocol": "aes128"
  },
  "switch_address": "10.23.34.45"
}
```

Response

Status: 201, Created

Name	Type	Description
authentication_method	string	Optional authentication method.
comment	string	Optional comment text.
engine_id	string	Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier.
name	string	SNMP user name.
owner	owner	Optional name and UUID of owning Storage Virtual Machine (SVM).
snmpv3	snmpv3	Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
switch_address	string	<p>Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.</p> <ul style="list-style-type: none"> • example: 10.23.34.45 • Introduced in: 9.7 • readCreate: 1

Example response

```
{
  "authentication_method": "usm",
  "comment": "This is a comment.",
  "engine_id": "80000315055415ab26d4aae811ac4d005056bb792e",
  "name": "snmpv3user2",
  "owner": {
    "name": "svml",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "snmpv3": {
    "authentication_password": "humTdumt*@t0nAwa1l",
    "authentication_protocol": "sha2_256",
    "privacy_password": "p@**GOandCLCt*200",
    "privacy_protocol": "aes128"
  },
  "switch_address": "10.23.34.45"
}
```

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2621475	This operation is not allowed on a node SVM.
2621601	This operation is not permitted on a system SVM.
5636101	One role for all applications and authentication methods is recommended for a user account.
5636106	Engine ID starting with first bit set as 1 in first octet should have a minimum of 5 or maximum of 32 octets.
5636107	Engine ID specified is based on an IP address, and must therefore be 9 octets in length.
5636108	Engine ID specified is based on an IPv6 address, and must therefore be 21 octets in length.
5636109	Engine ID specified is based on a MAC address, and must therefore be 11 octets in length.
5636110	Engine ID starting with first bit set as 0 in the first octet needs to be 12 octets in length.
5636111	Engine ID cannot have all zeroes.
5636112	Invalid Engine ID specified.
5636113	Cannot enable SNMP user capability for data SVM.
5636116	Cannot read SNMP user capability for data SVM.
5636121	User account name is reserved for use by the system.
5636123	Cannot create an SNMP user with a role other than readonly, none, or admin.
5636124	Cannot create an SNMP user with a role other than vsadmin-readonly, none, or vsadmin.
5636126	Cannot create a user with the user name, group name, or role "autosupport" because it is reserved by the system.
5636148	Failed to discover SNMPv3 engine ID of remote SNMP agent.
5636150	The switch_address parameter can be specified only for administrative SVMs.
5636152	0.0.0.0 cannot be specified as the IP Address for the switch_address parameter.
5636153	Engine ID is already associated with user.
5832711	The only application permitted for user "diag" is "console."
7077897	Invalid character in username. Valid characters:" A-Z, a-z, 0-9, ".", "_", and "-". Note that the user name cannot start with "-".

Error Code	Description
7077898	The username must contain both letters and numbers.
7077899	Username does not meet length requirements.
7077906	Cannot use given role with this SVM because a role with that name has not been defined for the SVM.
9043995	Failed to create SNMP community. Reason": SNMPv1 and SNMPv2c are not supported when cluster FIPS mode is enabled.
9043999	ONTAP failed to create an SNMPv3 user because SNMPv3 is disabled in the cluster.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

owner

Optional name and UUID of owning Storage Virtual Machine (SVM).

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snmpv3

Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
authentication_password	string	Authentication protocol password.
authentication_protocol	string	Authentication protocol.
privacy_password	string	Privacy protocol password.
privacy_protocol	string	Privacy protocol.

snmp_user

An SNMP user can be an SNMPv1/SNMPv2c user or an SNMPv3 user. SNMPv1/SNMPv2c user is also called a "community" user. An SNMPv3 user, also called a User-based Security Model (USM) user, can be a local SNMPv3 user or a remote SNMPv3 user. A local SNMPv3 user can be used for querying ONTAP SNMP server over SNMPv3 and/or for sending SNMPv3 traps. The local SNMPv3 user used for sending SNMPv3 traps must be configured with the same authentication and privacy credentials on the traphost receiver as well. A remote SNMPv3 user is also configured on a remote switch and used by ONTAP SNMP client functionality to query the remote switch over SNMPv3. An SNMP user is scoped to its owning Storage Virtual Machine (SVM). Owning SVM could be a data SVM or the administrative SVM.

Name	Type	Description
authentication_method	string	Optional authentication method.

Name	Type	Description
comment	string	Optional comment text.
engine_id	string	Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier.
name	string	SNMP user name.
owner	owner	Optional name and UUID of owning Storage Virtual Machine (SVM).
snmpv3	snmpv3	Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
switch_address	string	<p>Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.</p> <ul style="list-style-type: none"> • example: 10.23.34.45 • Introduced in: 9.7 • readCreate: 1

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Manage individual SNMP users

Support SNMP users engine_id name endpoint overview

Overview

Defines, retrieves, updates and deletes an individual SNMP user.

Examples

Retrieves the details of an SNMP user

```
# The API:
GET "/api/support/snmp/users/{engine_id}/{name}"

# The call:
curl -H "accept: application/hal+json" -X GET "https://<mgmt-
ip>/api/support/snmp/users/80000315056622e52625a9e911a981005056bb1dcb/snmp
v1user2"

# The response:
{
  "engine_id": "80000315056622e52625a9e911a981005056bb1dcb",
  "name": "snmpv1user2",
  "scope": "cluster",
  "owner": {
    "name": "cluster-1",
    "uuid": "26e52266-a925-11e9-a981-005056bb1dcb"
  },
  "authentication_method": "community",
  "_links": {
    "self": {
      "href":
"/api/support/snmp/users/80000315056622e52625a9e911a981005056bb1dcb/snmpv1
user2"
    }
  }
}
```

Updates the comment parameter for an individual SNMP user

```
# The API:
PATCH "/api/support/snmp/users/{engine_id}/{name}"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
PATCH "https://<mgmt-
ip>/api/support/snmp/users/8000031505b67667a26975e9118a480050568e6f74/publ
ic" -d '{"comment":"Default SNMP community"}'

# The response:
200 OK
{}
```

Deletes an individual SNMP user in the cluster

```
# The API:
DELETE "/api/support/snmp/users/{engine_id}/{name}"

# The call:
curl -H "accept: application/json" -H "Content-Type: application/json" -X
DELETE "https://<mgmt-
ip>/api/support/snmp/users/8000031505b67667a26975e9118a480050568e6f74/snmp
user"

# The response:
200 OK
{}
```

Delete an SNMP user

```
DELETE /support/snmp/users/{engine_id}/{name}
```

Introduced In: 9.7

Deletes an SNMP user. The engine ID can be the engine ID of the administrative SVM or a data SVM. It can also be the SNMPv3 engine ID of a remote switch.

Related ONTAP commands

- `security login delete`
- `system snmp community delete`

Learn more

- [DOC /support/snmp/users/{engine_id}/{name}](#)

Parameters

Name	Type	In	Required	Description
engine_id	string	path	True	Engine ID of owning SVM or remote switch.
name	string	path	True	SNMP user name.
fields	array[string]	query	False	Specify the fields to return.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none">• Default value: 1• Max value: 120• Min value: 0

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2621475	This operation is not allowed on a node SVM.
2621699	This operation is not allowed on a system SVM.
5636123	Cannot create an SNMP user with a role other than readonly, none, or admin.
5636124	Cannot create an SNMP user with a role other than vsadmin-readonly, none, or vsadmin.
5832712	Cannot modify attributes for user "diag."
7077906	Cannot use given role with this SVM because a role with that name has not been defined for the SVM.
9043999	ONTAP failed to create an SNMPv3 user because SNMPv3 is disabled in the cluster.

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Retrieve SNMP user details

GET /support/snmp/users/{engine_id}/{name}

Introduced In: 9.7

Retrieves the details of an SNMP user. The engine ID can be the engine ID of the administrative SVM or a data SVM. It can also be the SNMPv3 engine ID of a remote switch.

Related ONTAP commands

- `security snmpusers -vserver <SVM Name> -username <User Name>`
- `security login show -application snmp -vserver <SVM Name> -user-or-group-name <User Name>`

Learn more

- [DOC /support/snmp/users/{engine_id}/{name}](#)

Parameters

Name	Type	In	Required	Description
engine_id	string	path	True	Engine ID of owning SVM or remote switch.
name	string	path	True	SNMP user name.
fields	array[string]	query	False	Specify the fields to return.
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When iterating over a collection, the default is 15 seconds. ONTAP returns earlier if either max records or the end of the collection is reached.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Response

Status: 200, Ok

Name	Type	Description
_links	_links	
authentication_method	string	Optional authentication method.
comment	string	Optional comment text.

Name	Type	Description
engine_id	string	Optional SNMPv3 engine identifier. For a local SNMP user belonging to the administrative Storage Virtual Machine (SVM), the default value of this parameter is the SNMPv3 engine identifier for the administrative SVM. For a local SNMP user belonging to a data SVM, the default value of this parameter is the SNMPv3 engine identifier for that data SVM. For an SNMPv1/SNMPv2c community, this parameter should not be specified in "POST" method. For a remote switch SNMPv3 user, this parameter specifies the SNMPv3 engine identifier for the remote switch. This parameter can also optionally specify a custom engine identifier.
name	string	SNMP user name.
owner	owner	Optional name and UUID of owning Storage Virtual Machine (SVM).
scope	string	Set to "svm" for interfaces owned by an SVM. Otherwise, set to "cluster".
snmpv3	snmpv3	Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
switch_address	string	<p>Optional remote switch address. It can be an IPv4 address or an IPv6 address. A remote switch can be queried over SNMPv3 using ONTAP SNMP client functionality. Querying such a switch requires an SNMPv3 user (remote switch user) to be configured on the switch. Since ONTAP requires remote switch user's SNMPv3 credentials (to query it), this user must be configured in ONTAP as well. This parameter is specified when configuring such a user.</p> <ul style="list-style-type: none"> • example: 10.23.34.45 • Introduced in: 9.7 • readCreate: 1

Example response

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  },
  "authentication_method": "usm",
  "comment": "This is a comment.",
  "engine_id": "80000315055415ab26d4aae811ac4d005056bb792e",
  "name": "snmpv3user2",
  "owner": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  },
  "scope": "svm",
  "snmpv3": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "authentication_protocol": "sha2_256",
    "privacy_protocol": "aes128"
  },
  "switch_address": "10.23.34.45"
}
```

Error

Status: Default, Error

Name	Type	Description
error	error	

Example error

```
{
  "error": {
    "arguments": [
      {
        "code": "string",
        "message": "string"
      }
    ],
    "code": "4",
    "message": "entry doesn't exist",
    "target": "uuid"
  }
}
```

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

Name	Type	Description
self	href	

owner

Optional name and UUID of owning Storage Virtual Machine (SVM).

Name	Type	Description
_links	_links	
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snmpv3

Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
_links	_links	
authentication_protocol	string	Authentication protocol.
privacy_protocol	string	Privacy protocol.

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Update the comment parameter of an SNMP user

PATCH /support/snmp/users/{engine_id}/{name}

Introduced In: 9.7

Updates the comment parameter of an SNMP user.

Optional properties

- `comment` - Comment text.

Related ONTAP commands

- `security login modify`

Learn more

- [DOC /support/snmp/users/{engine_id}/{name}](#)

Parameters

Name	Type	In	Required	Description
engine_id	string	path	True	Engine ID of owning SVM or remote switch.
name	string	path	True	SNMP user name.

Name	Type	In	Required	Description
return_timeout	integer	query	False	<p>The number of seconds to allow the call to execute before returning. When doing a POST, PATCH, or DELETE operation on a single record, the default is 0 seconds. This means that if an asynchronous operation is started, the server immediately returns HTTP code 202 (Accepted) along with a link to the job. If a non-zero value is specified for POST, PATCH, or DELETE operations, ONTAP waits that length of time to see if the job completes so it can return something other than 202.</p> <ul style="list-style-type: none"> • Default value: 1 • Max value: 120 • Min value: 0

Request Body

Name	Type	Description
comment	string	Optional comment text.
owner	owner	Optional name and UUID of owning Storage Virtual Machine (SVM).

Example request

```
{
  "comment": "This is a comment.",
  "owner": {
    "name": "svm1",
    "uuid": "02c9e252-41be-11e9-81d5-00a0986138f7"
  }
}
```

Response

Status: 200, Ok

Error

Status: Default

ONTAP Error Response Codes

Error Code	Description
2621475	This operation is not allowed on a node SVM.
2621699	This operation is not allowed on a system SVM.
5636123	Cannot create an SNMP user with a role other than readonly, none, or admin.
5636124	Cannot create an SNMP user with a role other than vsadmin-readonly, none, or vsadmin.
5832712	Cannot modify attributes for user "diag."
7077906	Cannot use given role with this SVM because a role with that name has not been defined for the SVM.
9043999	ONTAP failed to create an SNMPv3 user because SNMPv3 is disabled in the cluster.

Definitions

See Definitions

href

Name	Type	Description
href	string	

_links

owner

Optional name and UUID of owning Storage Virtual Machine (SVM).

Name	Type	Description
name	string	The name of the SVM.
uuid	string	The unique identifier of the SVM.

snmpv3

Optional parameter that can be specified only for an SNMPv3 user i.e. when 'authentication_method' is either 'usm' or 'both'. This parameter defines the SNMPv3 credentials for an SNMPv3 user.

Name	Type	Description
authentication_protocol	string	Authentication protocol.
privacy_protocol	string	Privacy protocol.

snmp_user

An SNMP user can be an SNMPv1/SNMPv2c user or an SNMPv3 user. SNMPv1/SNMPv2c user is also called a "community" user. An SNMPv3 user, also called a User-based Security Model (USM) user, can be a local SNMPv3 user or a remote SNMPv3 user. A local SNMPv3 user can be used for querying ONTAP SNMP server over SNMPv3 and/or for sending SNMPv3 traps. The local SNMPv3 user used for sending SNMPv3 traps must be configured with the same authentication and privacy credentials on the traphost receiver as well. A remote SNMPv3 user is also configured on a remote switch and used by ONTAP SNMP client functionality to query the remote switch over SNMPv3. An SNMP user is scoped to its owning Storage Virtual Machine (SVM). Owning SVM could be a data SVM or the administrative SVM.

Name	Type	Description
comment	string	Optional comment text.
owner	owner	Optional name and UUID of owning Storage Virtual Machine (SVM).

error_arguments

Name	Type	Description
code	string	Argument code
message	string	Message argument

error

Name	Type	Description
arguments	array[error_arguments]	Message arguments
code	string	Error code
message	string	Error message
target	string	The target parameter that caused the error.

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.